INTRODUCTION TO THE C PROGRAMMING LANGUAGE

See Pages 20-38

Also in this Issue

and more

Complete Table of Contents on Page 3

Staff

In This Issue

Departments

NEWS & VIEWS

by Sol Libes

Digital Research Releases MP/M-II

Digital Research has introduced MP/M-II, which is specifically designed for multi-user systems. MP/M-II features file and record locking, password protection and increased disk capacity. The first version of MP/M, which was released in 1979, was well suited for multitasking applications but lacked provision for preventing concurrent access of a data record by more than one user.

Altos Computer Systems was the first to sign a license agreement for MP/M-II for its four-user system. Altos plans to enhance their version with interrupt-driven console I/O and other features to improve performance. Digital Research claims that as many as fifteen software firms have declared their intentions to extend their software to run under MP/M-II.

MP/M-II will run on an 8-bit system, with 48K RAM and clocked interrupt. It can support up to sixteen users, sixteen printers, sixteen disk drives (512Mbytes each max., 8Gbytes max.) and 400Kbytes of RAM. RAM can be divided among eight users in 48K banks. The MP/M-II package includes RMAC (relocatable macroassembler), LINK-80 (linker) and LIB (program library management utility). It is compatible with CP/M files (up to 32Mbytes max.) and supports CP/NET, allowing interconnection of multiple systems. MP/M-II is priced at $450 and OEM discounts are available.

Digital Research Publishes Software Catalog

Digital Research has published a *CP/M Compatible Software Catalog* which lists 116 independent software vendors (ISV) who write CP/M compatible software. The 24 page catalog is $5, and is available from: Digital Research, Box 519, Pacific Grove, CA 93950. ISV's who would like to be listed in the catalog should write to Digital Research asking for the appropriate forms to be sent.

Digital Research is rumored to be working on software based on the Motorola 68000. Most likely the first package to be released will be MP/M— look for it to be released in the last quarter of next year. Further, you can expect enhancements to MP/M-II to be released soon. Also due soon from the language group (formerly Compiler Systems) will be a true Basic Compiler producing object code that runs without a run-time program. An 8080 version is due in early 1982 with an 8086 version to follow in mid- or late 1982. The Basic Compiler will have several enhancements, including labels for writing structured-type programs. The group also plans other languages in addition to the present Basic and PL/I compilers. It is also rumored that CP/M 3.0, originally expected by year-end, has been put on a back burner.

North Star Unwraps Advantage & New Horizons

North Star has introduced several new computers: the "Advantage," a $4,000 desktop computer with graphics ability and two new multi-user systems based on the Horizon mainframe. The Advantage uses a Z80A with 64K-RAM plus 20K-RAM graphics memory (240 x640 pixels). It also features two double-sided, double-density mini drives, an 87-key keyboard (fifteen programmable keys) and a 12" screen. The Horizon systems support up to five users, and use either a 5Mbyte or 18Mbyte winchesters. Local networking interfaces are planned for all three systems.

Typesetting From CP/M ASCII Files

In one of my recent editorials I had mentioned sending a CP/M text file over the phone line to a typesetter, and having type set directly. Several readers wrote to me asking for the name and address of such a typesetter. I am sure that you can locate one in most of the larger metropolitan areas (just look in the yellow pages). If you cannot find one in your area, then I recommend you try Dale Brown, Brown Graphic Press, 2488 Summit St., Columbus, Ohio, 43202.

You will have to add typesetting control codes to the text that specify type size, spacing, style, etc. You must pay an initial fee of $50 to cover the cost of a manual and the proofing and correction of your first few submittals. Turnaround time is about two working days (including proofing), and about one day when no proofing is required. They can accept text on TRS-80 Scripsit or CP/M 8" disks. The modem service operates 24 hrs/day and you can save money by sending the file after 11PM.

Marc Author Killed in Accident

Ed Ziemba, author of the Marc Operating System (see the Jul/Aug 1981 *Microsystems*, page 12) died in June in a snorkeling accident. Despite this tragic loss, Leor Zolman, who was working with Ed on Marc, said that he and his associates plan to do the final polishes of the software and will release it shortly.

4   MICROSYSTEMS

Cambridge Development Lab Offers Newsletter And Seminar

Cambridge Development Laboratory is offering free subsrcriptions to its "Dynamic Blackboard News," a newsletter covering applications, hints and technical notes for graphics users. CDL will also present a one-day seminar titled "Introduction To Microcomputer Graphics" in Boston on November 14th ($20/person). To obtain more infor-

mation on the seminar, or to receive the newsletter contact: Jean L. Graef, CDL, 36 Pleasant St., Watertown, MA 02172; (617)926-0869.

User Groups News

The SIG/M group (Box 97, Iselin, NJ 08830) has released seventeen new volumes, bringing their total to 42 volumes. Thirteen more new volumes are in development, and are expected to be released by year-end. The CPMUG group (1651 Third Avenue, New York, NY 10028) has released four new volumes, bringing their total to 54 volumes. A full report on these two user groups will be published in the next issue of MICROSYSTEMS.

The Connecticut CP/M User Group meets on the last Monday of each month at 7:30PM at the Wordsmith Network, 110 Day Hill Rd., Windsor, CT 06095. For more information call Charlie Lowderback at (203)683-2427. An attendance fee of $2/meeting is charged.

"TCS Debits & Credits" is a new newsletter for users of TCS accounting systems software. It contains new product and version announcements, bug reports, corrections, and enhancements on TCS products. It will be published quarterly, and a subscription is $20/yr. For more information contact: TCS System Users Group, c/o Mountain Software Systems, Box 3282, Walnut Creek, CA 94598, (415)625-1592.

The Amethyst User Group has been formed to support the Mince, Scribble, and Amethyst text editor and formatter. The main purpose of the group is to provide coordination among users developing extensions to Mince and Scribble. Membership is $6/yr. For more information contact Barry A. Dobyns, 1633 Royal Crest #1128, Austin, TX 78741, (512)441-9466.

The Osborne Business Software Users Group has been formed at: 2256 Main St., Suite 11, Otay, CA 92011; (714)423-0538. Membership is $10/yr domestic, $20/yr foreign.

I recently received issue #8 of the Pascal/Z User Group Newsletter. The issue was 31 pages long, and it keeps improving. The group supports Inter-Systems Pascal/Z and has also just released Volume #11 of their public domain software. For information write: Pascal/Z UG, 7962 Center Parkway, Sacramento, CA 95823.

An Intertec Users Group has been formed called "Super*Star." It will pub-

lish a monthly newsletter and furnish other services. For information, contact: Amy Roosevelt, 3722 Chestnut Pl., Denver, CO 80216, (303)623-7973.

Digital Research has formed a PL/I-80 User Group (PLUG). PLUG members will receive the PLUG and ISV newsletters (Independent Software Vendor). A PLUG library has also been formed, consisting of source code programs submitted by PL/I-80 users. Two disks are already available to members at $25 each. Users who submit acceptable programs to the library or attend an ISV seminar will receive a year's free membership. For more information write: PLUG, Digital Research, Box 579, Pacific Grove, CA 93950.

The BDS-C Users Group (Box 287, Yates Center, Kansas 66783) has released three new volumes of software. One volume contains an updated version of Small-C (in source) which includes an assembler, linker, submit and clock board driver programs. The other two volumes contain utilities and a modem program.

New Bits

A 10-page quick reference guide to Oasis Basic is available for $1.00 from Phase One Systems Inc., 7700 Edgewater Dr., Suite 830, Oakland, CA 94621...........The second edition of the "CP/M Software Index" has been published listing 740 programs offered by 248 software vendors ($6 North America, foreign $8, CA residents add 6% tax). Order from: Small Systems Group, Box 5429, Santa Monica, CA 90405, (213)392-1234......Lifeboat Associates, the world's largest publisher/distributor of CP/M-based software has announced that it will soon begin distributing software for the new IBM personal computer DOS developed by Microsoft.

Forth Interest Group National Convention

The FORTH Interest Group (FIG) will hold a National Convention November 28, 1981, 9am - 6pm, at the Santa Clara, California Marriot Hotel. The one day convention will include a presentation, workshops, hands-on equipment and a number of vendor exhibits. An evening dinner will follow the day's activities and will feature a speaker. Registration is $3.00, dinner is extra.

For information: FORTH Interest Group, P.O. Box 1105, San Carlos, CA 94070, (415)962-8653. ∎

by Sol Libes

EDITOR'S PAGE

Another Language Issue

This issue of *Microsystems* continues to emphasize languages. We had originally planned to have only one issue on this topic (Sept/Oct 1981) and to emphasize MP/M in this issue. However, we received so many articles on languages that we had to change our plans. In fact, we have received enough excellent articles to have a third issue on languages—but we will hold off on that one for awhile. We have therefore set the following tentative schedule for *Microsystems*. Although we have some articles on these topics already, we are searching for more good material. I would appreciate hearing from potential authors on these topics:

Jan/Feb 1982 PL/I-80
Mar/Apr 1982 Multi-User Systems
May/Jun 1982 Data Base Systems
Jul/Aug 1982 68000 & Z8000 Systems
Sep/Oct 1982 Languages III
Nov/Dec 1982 Local Networks

IEEE-696/S-100
Bus Standard Status Report

I regret to report that the final revision to the IEEE-696/S-100 Bus Standard was not entirely ready as we were putting this issue together, early in September. We expect to print the entire standard as soon as it becomes available.

To print Or Not To Print,
That Is The Question

As editor of *Microsystems*, I have on a few occasions rejected an article that I really would have liked to publish.

The problem was that the article included source code that would have taken up an entire issue, and sometimes more than an entire issue. It was with deep regret that I rejected these articles.

Although the magazine has nearly doubled in size from the first issue published less than two years ago, editorial space in the magazine is still at a premium. We have only a certain number of pages allocated in each issue for articles.

Therefore, I am seriously considering omitting source code listings from the magazine when such listings are more than two or three pages long. Instead, we would refer the reader to the author to obtain the source code on disk. After all, who really wants to key in a lot of code with its attendant entry problems? Personally, I certainly would be willing to pay a reasonable sum to save the time and entry problems. Also, it would free up editorial space in the magazine so that we could include more articles.

Before making such a radical change, I would like to hear from Microsystems readers and authors. Are you in favor, or opposed to this change? What do you feel is a reasonable charge for the source code? Please let me hear from you at: Box 1192, Mountainside, NJ 07092.

IEEE-696/S-100
Interfacing Book Published

Mark Garetz and I have written a new book entitled *Interfacing To S-100/IEEE-696 Microcomputers* which has been published by Osborne/McGraw-Hill ($15.00). Naturally, it is impossible for me to discuss this book without prejudice since I am one of the authors. I feel however, that it is authoritative and comprehensive. It presents in-depth discussions of the S-100 bus signals, timing relationships, memory and I/O interfacing, interfacing to "real world" devices, digital-to-analog and analog-to-digital conversion, S-100 interrupt handling, using programmable timer/counters and setting up multi-master systems under the IEEE-696 TMA protocol. All circuit diagrams are complete with part numbers, values and pin numbers so that the circuits can be wired directly from the book. Several dozen software driver routines are also presented in the 321 page book for use with the circuits.

Digital Research
Makes Major Changes

Digital Research is undergoing major changes in organization. The company is formalizing its management structure, accepting venture capital financing, altering its pricing policies and moving to a new location.

Gary Kildall created CP/M in 1974 when working as a software consultant to Intel. Intel decided not to use it, so Gary formed Digital Research Inc. in 1975 to market CP/M. IMS took a look at it and decided to use it as the operating system for their Imsai-8080 microcomputer. It was an immediate success.

Editor's Page, continued...

In six years Digital Research has improved CP/M, added MP/M, PL/I-80, CP/M-86 and a number of software development packages. It has grown from a two person company (Gary and his wife Dorothy McEwan) to fifty employees. Sales for the fiscal year ending August 1981 were about $6 million.

Digital Research is now gearing up for what they expect to be a $10 million year and $100 million within next five years. First, they have given up a small share of ownership to four venture capital firms. The infusion of capital has allowed Digital Research to purchase Compiler Systems, developers of the CBasic compiler. Second, the money will make it possible for Digital Research to move out of the Kildall home and into a 16,000 square foot modern office building.

Digital Research has also changed the structure of its sales agreements. Past CP/M licenses required a one-time fee of $50,000 to microcomputer companies who wished to furnish CP/M. This proved to be a bargain for these vendors. Now Digital Research requires graduated payment based on the number of computers sold by the manufacturer. Digital Research claims to have licenses with about three hundred computer manufacturers. ■

Microsystems' Mix-up

Dear Editor:

A slight error crept into my article which appeared on page 42 of the March/April issue of *Microsystems*.

Those who made it to the third column (on page 43) must have been bewildered to find the sample code table (CODES1) where the code conversion instructions should have appeared. And, lo, where CODES1 belonged were the code conversion instructions. That section of the article should have read:

with the code character in the 'A' register:

```
D64A    SUI     FIRSTCH
FE0D    CPI     LIMIT + 1
D2nnnn  JNC     Error Routine
```

performs the conversion and validates the result. This leaves the offset in the 'A' register, so SELECTOR can be called as soon as the 'HL' register pair is set to the appropriate table.

If a code sequence isn't practical, the assembler will create a code table:

```
CODES1  DB  'E'  Code to be converted to zero
        DB  'P'  Code to be converted to 1
        DB  'T'  Code to be converted to 2
        DB  'D'  Code to be converted to 3
```

In this case....(etc).
Fred Gohlke
Carteret, NJ

The staff of Microsystems would like to extend apologies for this error to Mr. Gohlke and to the readers.

The CP/M Connection, Part IV

Dear Editor:

Congratulations on getting a much-needed magazine off the ground and into circulation! Articles on CP/M, such as Chris Terry's "The CP/M Connection," are desperately needed out here in the field, since the Digital Research documentation is often obscure, and books such as Dr. Zaks' recent contribution are written for a rank beginner, in addition to being poorly organized, hastily thrown together, and full of errors.

In Part IV of Mr. Terry's series, there is a discussion on finding the available memory size of a CP/M system. He quite correctly points out that the bottom memory location used by CP/M is contained in the address portion of the jump instructions at locations 0005H-0007H, where it may be read by the user's application program.

There is, however, one small "gotcha" that should be pointed out. One cannot simply crank up DDT and examine these locations to determine the bottom of CP/M, as DDT will change this address to reflect the reduced memory size in the debug mode.

Hence, while an application program can read this location and determine the top of available memory "on the run," you cannot get the correct answer by examining these locations manually with a debugger.
David Goodman
Marshall, VA

IDS-88/Cromemco CPU Fix

Dear Editor

Your article, "S-100 Modems," in the May/June issue gave me the clue I was looking for. I have been trying to get one of the IDS-88 modem boards to work with a Cromemco CPU—with no luck, as mentioned in the article. Spurred on by your note about a fix, I called IDS—about forty times over the last month. I wrote to them and, in desperation, even left a message on their modem phone (one of the most finicky I have ever seen, by the way!). No one is home—ever! One wonders if they are still in business, or how long they will be in business with that level of customer support.

Anyway, with that apparent dead end, I thought I would write to you. Could you please tell me what the fix is, or refer me to someone (a dealer?) who might know? I would appreciate any help you could provide.

While I'm writing—I like your magazine. It fills a void with a level of solid technical detail that is of much interest. Keep up the good work!
Larry Wright
2071 Oakley Ave.
Menlo Park, CA 94025

Editor's Note: *International Data Systems has filed for Chapter XI, and there is some doubt that they will be back in business. We never did discover exactly what the problem was between the IDS-88 Modem and Cromemco CPU cards. If any readers know, would they please let Larry (and Microsystems) know.*

COMMX Version 7

Dear Editor:

COMMX version 7 has been enhanced since the *Microsystems* product review by Glenn Hart (May/June 1981). The features added are:
1) IBM file transfer capabilities with TSO + Wilber.
2) The sending of "break" signal for half-duplex systems via control-B.
3) User dial directory for auto-dial PMMI or HAYES modem card types.

4) Remote control for use with "BYE" or "REMOTE" (PMMI or HAYES) programs.

I feel the following items need clarification regarding software operations and procedures that were misleading in the article:

All versions of COMMX are written entirely in 8080 assembler, allowing a purchaser of the source program to make modifications if desired.

The cost of the COMMX source program ($250) is *not* an obvious attempt to discourage its purchase, as Mr. Hart prematurely concludes. It is directly based on efforts to provide an extensively tested program along with immediate user technical support where professional experience is required. Since every line of source contains a comment, many useful programming techniques will be discovered in the twenty different CP/M function calls implemented. Both async- and synchronous file transfer protocol as well as other extended features are provided to suit the needs of any user. A file transmit option (not mentioned in the article) allows full or no echo wait with a connected host system. Full echo wait indicates echo error, providing a measure of transmit data integrity without any special requirements on the connected host system other than the ability to echo the characters it receives (termed echoplex). The operator can either abort the transfer or note the error location and continue the transfer.

COMMX synchronous protocol is different in concept from others and has major differences! (The article was very misleading on this point.)

• protocol optimization of overhead for reduced online time.

• error-free file transfers using CRC16 polynomial.

• fast CRC16 encode/decode (all registers in 43 bytes!).

• CRC16 formatted for transmission on timeshare ASCII line.

• real-time master/slave operation (XON/XOFF not used here).

• one byte controls consisting of STX, ETX,ACK,NAK,ENQ.

• allows transfer of any data type (COM, REL,INT,BAS,etc.).

COMMX is menu-driven, which means all commands are displayed on the screen and, all input options are prompted, making program operation child's plays (with simple keystrokes). Since most features are automatically selected by the program, it appears to have a reduced command set. Is a program with a larger command set easier to use? I suggest a closer look into the command set options before evaluating all the feature bits in a program's bucket!

COMMX (now version 7) functional documentation is available at no charge from Hawkeye Grafix, 23914 Mobile, Canoga Park, CA 91307.

New features with COMMX version 7 include: transmit break for IBM type half-duplex attention control, file transmit no echo wait line pause and selectable continue character on IBM type half-duplex links, automatic line initialization to 8-bit data no parity (from 7-bit even parity) and console echo switched off with .COM file transfers, duplicate file create warning allows intentional override, auto-dialing modem types accept user created dial directories, auto redial, and remote control capabilities. COMMX still executes in 5K bytes, allowing the rest of available system memory for disk buffer.

Will Pierce
Hawkeye Grafix
Canoga Park, CA

⟨ The CP/M* Bus ⟩

by Anthony Skjellum

The topic of discussion this month is once again enhancements for CP/M2. Due to space limitations, the section on the internals of link records will have to be deferred to the next installment.

Sub-directories

Operating systems on large computer systems usually provide sophisticated filing systems which allow sub-directory structures to be established within a user's main directory. This allows various files to be collected in a logical way. CP/M2 provides no such named sub-directories, although it provides the USER feature which allows files to be separated into sixteen categories on a given disk. Hard disk and high density floppy systems are making large amounts of on-line storage available on CP/M systems. However, greater storage capacity leads to problems in data organization which can be alleviated by the use of sub-directories.

At the outset of this series of enhancements for CP/M, we stated that a major premise of enhancements would be to stay within the original architecture of the software system. Therefore, in discussing sub-directories, we will propose a conservative approach.

The above consideration limits our choices for implementing sub-directories, but does not remove the possibility of making a powerful but feasible expansion to CP/M. To provide sub-directories, we will define a new variety of file to be known as a directory file. Directory files will have the file type of .DIR by default; but this will not be required. A status bit (denoted as b4') will be used to indicate that the file is a directory. Directory files will reside in any of the sixteen user areas, but will not nest. Such a file will contain a fixed number (e.g. 64) of 32 byte entries in the same way the CP/M maintains its main directory. To a CP/M system, it will appear outwardly as a single file and the directory file will allocate all the disk space of its constituents. Thus, it will be possible to treat it as a whole or as a set of files for the sake of various CP/M operations.

Several new BDOS commands and CCP functions will be needed to provide the sub-directory feature. We will begin by describing the new BDOS commands needed.

To begin with, we will need a new BDOS command which will create the directory file. The command will be known as MKDIR. It will set the directory status bit b4' to indicate that the file is a directory. Furthermore, the A register will be passed with the number of directory entries to set up for the directory, and sufficient storage will be allocated to contain these directory entries. Also, the first record of the directory file will contain information including the number of entries which the directory can contain. Here is an example of a typical BDOS call using MKDIR:

```
MKDIR <dir>  ; <dir> is an unambiguous file name
```

This example creates a ten entry directory from the file name stored in the FCB marked "fcb." As shown in the above example, the accumulator is returned with zero for a successful operation. The value 1 will be returned for a file exists condition, and -1 for a disk write error condition.

Now that we have defined the way to make sub-directories, there must also be a BDOS command to change CP/M's working directory from the main directory to a sub-directory. It must also be possible to change back to the main directory when necessary. This command will be denoted as CD (which stands for change directory, as in Unix). When calling CD, the DE register will point to an FCB with the file name of the directory to change over to. Upon return, the accumulator will contain zero if the operation was successful. A value of -1 will indicate that the file did not exist, and 1 indicates that the file was not a directory. Also, passing a value of zero in DE will cause CP/M to return to the main directory.

Once the working directory has been changed to a subdirectory, the search first and search next functions will examine the sub-directory and operation will be totally transparent to transient programs. Furthermore, when the simple link facility provided by LN is used within a sub-directory, it will refer back to the file in the current main directory and not in user area zero. Also, sub-directories in user area zero will be permitted to have simple link files.

The BDOS file delete command will also remove files from a sub-directory in the expected way. However, it will be very simple to delete the directories themselves when in the main directory. Therefore, caution would have to be exercised. Certainly the CCP ERA command would have to be changed to confirm any directory deletion.

Besides the above commands, it will also be necessary to generalize the rename capability of CP/M to allow renaming of files between sub-directories, and between a sub-directory and the main directory. This will be provided by a single BDOS command called MV, again using the Unix convention for naming. On entry, the DE pair will point to an FCB with the old and new names for the file as in the standard rename command (see p. 20 of *CP/M 2.0 Interface Guide*). However, both the accumulator (A) and the HL register pair will be used to pass further information. Calling with the accumulator set to zero will indicate that we are requesting a directory-to-directory rename. In this case, HL must point to an FCB with the name of the source directory followed by the destination directory in the second 16 bytes of the FCB. This format is just like that used for the file names themselves pointed to by DE. Finally, the accumulator is returned with the same error codes as for a standard rename.

If the accumulator is non-zero, the BDOS will expect a transfer between a sub-directory and the main directory. The HL register will point to an FCB containing the name of the sub-directory. If A is odd, (e.g. 00000011B), then the transfer is made to the main directory. Otherwise the transfer will be made to the sub-directory pointed to by HL.

Certain shorthands will also be allowed within the FCB. A ^A as the first character of the file name will indicate the current directory, while a ^B will indicate the main directory. Also, return current directory name command will be supported. This command will return a pointer to the directory name so that transients can determine the sub-directory they are in. The null string is returned for the main directory. This command will also be used by a new CCP function called PWD (again after Unix) which will display the name of the working directory. For the main directory, PWD will display nothing.

CP/M Bus, continued...

To complement the new BDOS commands are several other new CCP commands. The command MKDIR will make a directory and has the following syntax:

```
MVI     C,MKDIR  ; command is mkdir
LXI     D,FCB    ; point to fcb with name of file
MVI     A,10     ; 10 entries for this directory
CALL    BDOS     ; execute call
ORA     A        ; see if error...
JNZ     ERROR    ; yes...
```

Also, a change directory command (CD) will be provided. It will take two forms:

```
CD <dir> ; change to a directory
CD       ; change back to main directory
```

Similarly, the REN command will support renaming between sub-directories and the main directory. By specifying a directory name in square brackets ([]) before the file name, REN will understand that it is to do an inter-directory move. Specifying [] indicates the main directory. An example is presented here for clarity:

```
CD TEMP                      ; change to directory TEMP.DIR
REN TEST.TXT=[]TEST.TXT      ; get TEST.TXT from main
REN [TEMP2]TEST.TXT=TEST.TXT ; move TEST.TXT from
                             ; TEMP.DIR to TEMP2.DIR.
```

Finally, the DIR command will be generalized to list the contents of directory files when the second argument D is specified.

One further CCP command will be needed. This will be the PATH command. The PATH command will indicate to CP/M which sub-directories it should search for .COM files after looking in the current directory. A BDOS command will also be provided. The BDOS command will allow a transient to give CP/M the address of a set of "half"-FCB's containing the names of the path for the logged disk. This will remain in effect until a cold boot. A transient will also be able to pick up the path set before it was invoked. This command will return an address vector to a set of "half"-FCB's containing the information.

The CCP PATH command mentioned above takes two or more arguments. The first argument is either a plus sign, a minus sign or a disk name. Plus indicates that the following list of directories are to be added to the path list, while a minus sign removes the list from the path. An alternate form of the command reports the current path of a disk. This is done by providing a disk name instead of the plus or minus sign. Examples follow:

```
PATH+ ASM1 ASM2 TEXT1 ;add ASM1.DIR, ASM2.DIR
                      ; and TEXT1.DIR to path
PATH - TEXT4.DOC      ; remove TEXT4.DOC
                      ; directory from path.
PATH B:               ; report current path of B:
```

The path feature should further reduce the need for extra copies of standard programs and utilities.

Some additions to PIP are also in order. In following the CP/M2 convention for handling user areas, we will allow PIP to copy into the current directory, but not into other directories. This feature will be provided by generalizing the G option. Currently, G is followed by the user area number. We now add the option of adding a colon and the subdirectory name (terminated by Z) to this specification. Thus, we might do something like the following:

```
PIP TEST.FIL=TEST.FIL[G0:SUBDIR^Z] ; this will get TEST.FIL
                                   ; from the sub-directory
                                   ; SUBDIR.DIR in user area
                                   ; zero.
```

The topic for the next installment will be a discussion of the internals of link records used in complex link files. ∎

An Introduction To
The C Programming Language

by David A. Gewirtz

An introduction to the C language and reviews of four popular microcomputer implementations

C is the language that was developed after B. Although its name isn't very original, the C language is a highly versatile programming tool. C cannot be classified either as a "higher" or a "lower" level language. Probably the best classification would be a low-level, higher level language. It has also been called a systems implementation language.

C has many of the higher level control constructs, such as while loops, if-else conditionals, and block structuring. C also easily manipulates machine-related data types such as the byte with operations like bit-wise shift This combination makes C very useful for programming operating systems, languages, and utilities. It can also be easily used for other applications as well.

C, originally developed at Bell Laboratories, has evolved from the BCPL language created by Martin Richards, and the B language written by Ken Thompson for the first UNIX operating system. It shares many of the features of both BCPL and B, although where they concentrated almost exclusively on machine words, C has been expanded to include larger integers, characters, and in some implementations, floating point numbers. In addition, the syntax is different between BCPL and C. C also borrows many ideas, like the typing of variables, from other languages including ALGOL, Pascal, and PL/I. The most noted application of C was the development of the UNIX operating system for the PDP-11 computer.

Unlike Basic, C is a compiler, not an interpreter. This means that once compiled, programs run much faster because the statements in the program do not have to be retranslated into machine code each time they are encountered, which can become a very time consuming operation. In addition, the resulting object program that is run takes up much less space because the interpreter does not have to be resident in memory while the program is running.

C does not have the line numbers used in most versions of Basic. Instead, a C program is organized into manageable routines called functions. These functions, combined with the ability to make compound statements, make programs very straightforward. Wherever a simple statement can be used, a compound statement can be used there instead. For example, in the following "b=9" is a simple statement:

```
if (a==5) b=9;
```

Any group of statements, enclosed in brackets " { " and " } ", can be substituted for that simple statement:

```
if (a==5)
    {
    b=9;
    c=36;
    if (count<8) ++count;
    }
```

The operators "==" and "++" will be explained later.

These compound statements have a recursive definition. A compound statement is a statement containing other statements. Any of the statements inside a compound statement can also be a compound statement. This very powerful way of using program control statements allows programs to be "block-structured." I will not argue the relative merits of structuring, save that it makes programs easier to design, and clearer to debug and understand. This same block-structuring construct can be found in ALGOL and Pascal.

Functions are usually small, stand-alone subprogram segments of the main program. They are similar to the

David A. Gewirtz, 38-67 Taylor Rd., Fairlawn, NJ 07410.

Basic gosub statement in that when they are called, control jumps to them, and at the end they can return to where they came from. That is the only similarity. In Basic, all variables are global. If a routine to position the cursor on a terminal at X and Y coordinates was used in Basic, it might look like this:

```
330 REM Position Cursor at A and B
340 X=A
350 Y=B
360 GOSUB 5000

4999 REM Direct Cursor Addressing
5000 PRINT CHR$(27)+"="+CHR$(32+Y)
     +CHR$(32+X);
5010 RETURN
```

You would not be able to use X and Y for anything but that routine, and would have to assign the actual values A and B each time the routine had to be called. In addition, you would have to remember that 5000 was the cursoring routine. The same thing in C would be the call:

```
cursor(a,b);
```

which is much clearer to understand. The function would be:

```
cursor(x,y)
       {
       outchr(27);
       outchr('=');
       outchr(32+y);
       outchr(32+x);
       }
```

Outchr is another function which puts one byte to the screen. Once the cursor function is known to be working, it can be put aside and ignored, save to remember to use the function cursor(x,y) whenever it might be needed.

C makes very extensive use of the function feature. The C compiler can be relatively small in terms of reserved compiler operations, and all extraneous operations can be custom designed to fit a user's needs. In fact, C does not have any defined I/O, rather it consists of machine language functions to be called when needed. With functions, programs are very simple to modify. For example, if a new terminal were to be used, all that would be required would be a change in the function. The program remains the same.

Most C compilers come with a run-time library that consists of many previously compiled standardized functions such as outchr and cursor. The programmer simply has to remember the rules of what goes in and comes out of each function, but doesn't even have to include them in the actual program file. They simply must be "linked" to the compiled program before it can run.

The flow of control throughout C programs also helps to promote the readability and structure of the programs. C has many statements that evaluate conditions and control program flow accordingly.

if (expression) < statement > [else < statement >]

> If a given condition is true, do something (a statement which can always be a compound statement), else (if it is not true) do something else. If the else is omitted, flow will continue directly after the statement with the if.

while (expression) < statement >

> While a given condition is true, do something until that condition is no longer true.

do < statement > while (expression)

> Do something while a given condition is true. This differs from the while in that the condition is checked after running through it, not before.

for (expr1; expr2; expr3) < statement >

> This works somewhat like the Basic FOR statement. It controls looping through a statement by evaluating three expressions. The first expression is evaluated once, then as long as expr2 evaluates to something other than zero, the statement is executed and expr3 is evaluated. This could be seen as:

```
for (i=n ; i==j/3 ; ++i) <statement>
```

> which says "for i becomes n, until i equals j/3, do the statement, and increment i." The symbol "++" is an increment operator which will be examined later.

goto < identifier >

> Branch to the label specified by < identifier >.

switch (expression) < statement >

> This can be seen as a very powerful ON-GOTO statement. It evaluates the expression, and if it matches any of the cases in the statement, it will execute from that point. A break statement will return control outside of the block. For example:

```
printf("What is your choice?");
getchr(choice);
```

> This will get the choice for some set of options. Toupper is a function to convert lowercase alphabetic characters to their uppercase equivalent. If "B" was chosen, the function called reboot would be executed. When it finished, the break would be encountered which would return control outside the switch. If "Z" were chosen, the screen would be cleared, and if none of the options listed were chosen, the error message would be printed.

In addition to the statements above, C has a few very powerful operators. Unary operators are operators that do an operation to a single piece of data. The following is a list of the C unary operators. All of the operators return a value to some expression. For example, in Basic, -X does not mean negate X, but rather returns the negative of X to a specific expression.

```
*p - Pointer to nnn.
     The value is the value of the object currently
     pointed to by nnn.

&x - Address of x.
     Returns a pointer to x.

+x - States that x is positive.
     Returns the same value of x.

-x - Negative of x.
     Returns the negative of x.

++x - Increment x.
     The result is the new value of x.

--x - Decrement x.
     The result is the new value of x.

x++ - Increment x.
     The result is the original value of x.

x-- - Decrement x.
     The result is the original value of x.
```

```
x - Returns the ones complement of x.

!x - Not x.
     The result is 1 if x is 0, otherwise 0.

(type-name)x - coerce x to the type type-name.
     Returns the value obtained by converting the
     value of x to that type.

sizeof x - The result is an integer value equal
     to the size in bytes of x.

sizeof (type-name) - The result is an integer value
     equal to the size in bytes of an object of type
     type-name.
```

There are also quite a few binary operators available for C:

```
x*y - Multiply.
     The result is the product of x and y.

x/y - Divide.
     The result is the quotient of x divided by y.

x%y - Remainder.
     The result is the remainder of x/y.

x+y - Add.
     The result is the sum of x and y.

x-y - Subtract.
     The result is the difference of y from x.

x<<y - Shift left.
     The result is x shifted left y bits.

x>>y - Shift right.
     The result is x shifted right y bits.

x<y - Less than comparison.
     The result is 1 if true, 0 otherwise.

x>y - Greater than comparison.
     The result is 1 if true, 0 otherwise.

x<=y - Less than or equal comparison.
     The result is 1 if true, 0 otherwise.

x>=y - Greater than or equal comparison.
     The result is 1 if true, 0 otherwise.

x==y - Equal to comparison.
     The result is 1 if true, 0 otherwise.

x!=y - Not equal to comparison.
     The result is 1 if true, 0 otherwise.

x&y - And.
     The result is the bit-wise and of x and y.

x y - Exclusive or.
     The result is the bit-wise exclusive or of x
     and y. Be careful not to confuse this with the
     "power of" facility of Basic.

x y - Inclusive or.
     The result is the bit-wise inclusive or of x
     and y.

x&&y - Logical connective and.
     The resuslt is 0 if x is zero, without
     evaluating y. Otherwise the result is 1 only
     if both x and y are non-zero. This is used in
     logical expressions such as:

     if (a==b && b==c) a=c;

Evaluation stops as soon as the expression
     becomes false.

x y - Logical connective or.
     The result is 1 if either x or y is non-zero.
     This is used in expressions such as:

     while (a!=b  b!=c  d!=q) ++a;

t?/x:y - Conditional evaluation.
     This acts like the statement:
```

```
     if (t)
          (
          x;
          )
     else
          (
          y;
          )
```

```
The result is the result of the evaluation.

x=y - Assignment.
     The result (which is x) is the value of y.

x*=y - Equivalent to x = x * y.

x/=y - Equivalent to x = x / y.

x%=y - Equivalent to x = x % y.

x+=y - Equivalent to x = x + y.

x-=y - Equivalent to x = x - y.

x<<=y - Equivalent to x = x << y.

x>>=y - Equivalent to x = x >> y.

x&=y - Equivalent to x = x & y.

x =y - Equivalent to x = x   y.

x =y - Equivalent to x = x   y.

x,y  - x is evaluated first, then y.
     The result is the value of y after evaluation.
```

Although these operators make writing programs much easier and less verbose, their conciseness tends to make the programs difficult to read.

Variable and function names can generally be of any length, and use any displayable characters, as long as they do not conflict with the reserved words. Data in C is organized in both types and classes. A storage class is a method of storing the data. C has the following storage classes:

extern — Specifies that an external definition for the given identifier outside of that file.

auto — Specifies that the given identifier will exist as dynamic local variable. It will be created when that block of code is executed, and destroyed when the block is exited.

static — Specifies that the given identifier will not be known outside a block, but will remain inside that block. It differs from auto in that the data declared static will always exist until the termination of the program. This is similar to Algol's "own" declaration.

register — Means the same as auto, but specifies that efficient storage, such as registers be favored to hold the object, and the address of the object (&x) cannot be taken.

typedef — Specifies that the identifier should be recognized as a type specifier. This is used to create new data types from old ones. This is similar to Pascal's "type" declaration.

Data types are relatively primitive in C. Arrays are constructed of scalars; many implementations do not have them at all. The following is a list of data types that might be available:

char — A byte integer used to hold a single character of the machine's character set.

int — Usually a two byte integer, can also be called "short" in some versions.

long — Usually a four byte integer.

float — A floating point number, usually four bytes.

double — A double precision floating point number usually eight bytes.

struct — A sequence of one or more types used to create logical records in a program. Similar to Pascal's "record" and COBOL's record. In some implementations can be used to access individual bits in a byte as a field.

union — A record specified by fields.

pointer to — An unsigned integer used to hold an address of some object.

```
char line[80], *pline;
```

*pline is the pointer to declaration.

array of — A repetition of some type. Example is above in line 80 declares 80 repetitions of type char, and aligned it on line.

Pointers provide a very powerful and sometimes confusing tool in program writing. Simply stated, a pointer points to some location in the machine. They can always point to any location in the machine's address space. A pointer can be manipulated, or by placing an asterisk in front of the name, that which is being pointed to can be manipulated. For example:

```
char fcb[12], *pfcb;
```

Declares an array of 12 bytes with the first byte called fcb. If fcb is located at 54A2h, then the following would set the pointer pfcb to 54A2h:

```
pfcb=&fcb;
```

Note the value of pfcb has been changed to 54A2h. If we wanted to increment the pointer to point to the second byte, the following could then be done:

```
++pfcb;
```

Now, if we want to set what is now at 54A3h to "?", then the following could be done using indirection through pointers:

```
*pfcb='?';
```

This is where C can be confusing:

```
*++pfcb='?';
```

does the same thing. It increments pfcb to 54A3h and puts a "?" at that address.

C is not perfect however, there are some major disadvantages as well. One is the problem of operators being intermixed in confusing ways, as with *++pfcb.

Another rather serious problem is in the definition of the language itself. There is no actual formal definition of the C programming language. The closest is in a book written by the designers of the language [Kernighan 1978] in which there are quite a few vague points where some things are not absolutely explicit. Because of this, some implementations may differ in points of interpretation.

The next problem is that of portability. Although many implementations claim to be transportable to systems using compilers other than their own implementations, such as UNIX systems, they may in actuality be transportable to only a few other systems. This is not because of the syntax of the language, which will probably be identical, but because of the definition of the functions used. It may be necessary to rewrite some of the functions used to do what you expect them to do.

Lastly, and this is a mixed blessing, the C compiler is not overly picky on the whole. It will allow types to be confused without doing anything about it, and this can lead to difficult code to read, as well as results that may be CPU dependent. On the other hand, if you want to do something fancy on a specific machine, it will allow it.

Although there are drawbacks, C is a very powerful higher level language that is simple to use. It is one of the best for doing any form of systems software, and has quite a large user following.

Name: BD Software C Compiler
Price: $145
Distributed by:
Lifeboat Associates
1651 Third Avenue
New York, NY 10028

The BDS C Compiler was written by Leor Zolman of Cambridge, MA, and is distributed by Lifeboat Associates in New York City. This compiler takes the large step from a small development language to a full scale production compiler. Although not nearly as large or complex as compilers residing on mainframe computers, the BDS C compiler can handle a significant portion of the programming requirements of a microprocessor-based small computer.

Unlike Basic, it is not an interpreter but an actual compiler that generates machine code. Unlike most Pascals, it is a native code compiler. That is, it generates machine code directly executable on the machine for which it was designed (in this case the 8080, Z-80, and 8085 microprocessors) without generating an intermediate source code such as assembler source code or Pascal "P-code".

The BDS C package comes with a CP/M-compatible floppy disk containing a version of the compiler (in this case version 1.43) and the linker. In addition, all of the support modules containing compiled or assembled functions, as well as the run-time package, are also included. Although the compiler itself is not furnished in source form, the C run-time package and the two standard function libraries written in C, and the standard assembly language function libraries are furnished in ready-to-run versions and as source files. Since any C program is (or should be) made up of many separate functions, the BDS C package contains a library manager called CLIB that allows the programmer to build library files out of compiled functions from many files.

Although the BDS C Compiler does not directly support floating point numbers, a set of functions is included that allow floating point numbers to be manipulated in a reasonable manner when necessary. Another set of functions allows the programmer to integrate assembler code into callable functions. These routines require the assistance of the Digital Research Macro Assembler however, which is not included on the disk. Lastly, a

number of example programs are included, such as a working Othello and a rather nice terminal emulation program called "Telnet." Once Telnet has been customized for a particular computer, it provides intelligent terminal operations such as file transfer and receive. With another computer also running Telnet, it will transfer compiled files as well. I have run Telnet for quite some time to provide communication between my computer's acoustic coupler and a PDP-2060, as well as a PDP-1050, and Micronet.

In addition to the floppy disk software and documentation, BDS C comes with a seventy page manual that describes the BDS implementation, how it varies from previous BDS C versions, how to use the compiler, the linker, and the librarian. Also listed is usage documentation of the user available functions. Finally, comparision of BDS C to the "standard" C specified in Kernighan and Ritchie is also included.

Included in the BDS C package is a book by the authors of the original C language from Bell Labs, *The C Programming Language*, by Brian W. Kernighan and Dennis M. Ritchie. If anyone seriously intends to program the authoritative reference on C and is a necessary item in any C programmer's library.

In addition to the floppy disk software and documentation, BDS C comes with a seventy page manual that describes the BDS implementation, how it varies from previous BDS C versions, how to use the compiler, the linker, and the librarian.

BDS C complies with the syntax specified in Kernighan and Ritchie and supports many of the features of C—with some exceptions. These include short int, long int, float and double. Also there are no explicitly declared storage classes. Static and register classes do not exist; all others are either external or automatic. That is, if a function is called and it is not in that file, it is simply left as a reference to be resolved by the linker. Initializers are not allowed, so you cannot declare a variable and assign it a value in the declaration section. Unary operators "(type-name) expression" and "sizeof (type-name)" are not implemented. In structure and union declarations, bit fields are not implemented.

However, the preprocessor functions #define, #include, #ifdef, #ifndef, #else, #endif, and #undef are implemented.

The functions available with BDS C are reasonably complete with many functions from standard C, such as printf, getc, ungetc, scanf, and fopen. Simple machine-oriented functions that call the CP/M BDOS and the BIOS are available. Peek and poke memory locations, input and output from ports, call assembly language functions, seed and get random numbers, fill a block of memory, move a block of memory, sort a set of elements,

execute a CP/M program, and a primitive set of storage allocations functions are also included. Character input and output, with and without formatting is available. Many string and character processing functions are included, as well as file input/output functions. Lastly, a set of plotting functions for DMA video boards such as the Processor Tech VDM-1 are included (I have not tried them).

BDS C compiles in three stages. The first stage is the parser and the second is the code generator. The output of the code generator is a C Relocatable (CRL) file. This CRL file then has to be linked with all of the functions and such to generate the final runable file using CLINK. If no options are specified, CLINK produces a COM file that is executable by CP/M. There are options available in both the compiler and linker that allow programs to be executable at addresses other than 100 Hexadecimal as CP/M requires. This is a great advantage to those people who must write software for other systems or system software for CP/M. The BDS C compiler compiles in one chunk, loading the entire source file into memory at once as opposed to loading the source in segments, working on it a bit, and producing the CRL file. That means that the maximum size of a source file is limited by the size of main memory. This problem can be circumvented by writing programs with a lot of functions and then compiling the functions separately. I have done quite a bit of work with large C programs using many functions, and have had no problem using the compiler in 48K bytes of available memory.

The compiler is simple to use, and compilation and linkage is fast. I have written programs in BDS C and assembler of equal function and found the BDS C compiler to produce executable code in less time than CP/M's assembler. Unfortunately, the assembler produces smaller code than the compiler, since every BDS C program must include approximately 2K of run-time package with every program.

When I first tried to bring up version 1.32 of the compiler, I had no problem. I then received version 1.43 from Leor Zolman, and was unable to link any of the files that I had compiled. I called Leor up and told him of the problem. Since I had to go into Boston anyway, I decided to bring in my non-working version and pick up a working version. I stopped at his place where he gave me a disk that was supposed to be working. I went home to try it out, and when I did, it still didn't work (it proceeded to happily crash my system every time I tried). I called Leor and explained all of the symptoms. He had no idea what was wrong, but offered to make the forty plus mile trip to Worcester (which is horrible in itself) by motorcycle on a very cold day to see if he could figure out what was wrong. To make a long story shorter (I hope), he came out, we tried to figure out what was wrong, and it turned out I had a bad copy of the run-time package.

Other than that one incident, I have been impressed with the operation of the compiler. An interesting note about the BDS implementation is its user following. I have met many people who have been highly impressed by the compiler, and there is also a user group for it. In addition, both the Small C manual and Tom Gibson of Tiny C recommend the BDS C compiler as well. It is a complete and inexpensive implementation of a C compiler

MAGIC TYPEWRITER ("Word Processing and Database Management in one program") is an all-purpose tool that eliminates the need for separate programs for separate tasks. Word processing features include odd-even page headings and page numbers, full justification, centering, or right and left justification, and embedded commands to alter format specifications within text. As a database management system, Magic Typewriter allows the user to customize the system to his own needs. You could buy a lot of different software—a fast-sort, a mailing list program, a text-formatter, a database manager, and half a dozen other special programs—or you could simply buy Magic Typewriter, the one program that does it all. Magic Typewriter is available in North Star DOS and CP/M at a cost of only $175. The manual alone is $20. Include $2 postage and handling. We accept Visa and MasterCard, and California residents must add 6% sales tax. For further information, ask for our free brochure.

CALIFORNIA DIGITAL ENGINEERING
P.O. BOX 526 ★ HOLLYWOOD, CA 90028

Introduction to C, continued...

and I would recommend it to anyone intending to do any serious work in the C programming language.

```
Name: Small C Compiler
Price: $15
Distributed By:
     The Code Works
     P.O. Box 550
     Goleta, CA 93017
```

The Small C compiler first made its debut in an article published by Ron Cain in the May 1980 issue *Dr. Dobb's Journal*. He opened the article by saying that he had to have a compiler for his home computer. Since Cain did not have CP/M and did not wish to spend the money for the expensive compilers that were on the market, he decided to write his own. He chose to write his 8080 flavored C compiler in C, first using the Tiny c interpreter, then using a full UNIX C compiler. While writing it, he was careful to be sure that all of the code used in writing the compiler would potentially be able to be compiled by that compiler, which would allow it to be modified once installed in the target system. Upon completion of that project, he decided to publish the entire compiler, written in C, in *Dr. Dobbs Journal* along with a discussion of the operation of the compiler. For someone who really needed a compiler, this was great—that is if you had access to a machine with a C compiler on it and if you wanted to type in ten pages of (two columns each) of very, very small print.

At about this time, a company called The Code Works from Goleta, CA entered the scene. They took Ron Cain's Small-C compiler and interfaced it to CP/M. They added to it an unpublished (at that time) run-time system, a few sample programs, and an eleven page instruction manual. They put it all, minus the manual, on a standard CP/M eight inch floppy disk, which included all of the source code for everything, and a working compiled version of the compiler that generated code acceptable to CP/M's assembler. After all that, they started to market it, selling it for an unprecedented $15 plus two dollars postage.

When their press release came to *Microsystems* for publication, I thought something had to be in error. I hadn't seen any compiler software commercially sold for CP/M use at anywhere near that price. When I called them, I was assured that this was in fact the price of the package and not a typographical error.

Once I examined the Small-C package, I was impressed. There was none of the quick, throw-it-together programming and documentation I have found in some of the other software available at that price range. Instead I found a clear, concise manual describing what you could and could not do with this compiler, including a warning that it was not suited for major systems programming tasks. It explained the way its functions were used, and how to interface to assembly language. At the end, the manual gave a specification of what features were available from standard C.

Small C does in fact comply with the "standard" C syntax defined by Kernighan and Ritchie, however its function calls do not comply with standard nor does it support many of the nicer features of C. It has no storage class specifications. The compiler is limited to two byte integer and single byte character types, as well as pointers. In addition, single dimension vector arrays are available. Most unary operators are supported, with the exception of not "!", complement "^ ", and "sizeof". Most binary operators are supported as well, with the exception of connectives and "&&" and or " || ". This means that such statements as:

```
if (a<b && c>d) ++b;
```

cannot be used. The if-else abbreviation "?." is also not available, but is not really a great loss. The concatenated assignment operators "+=", "-=", "*=", "/=", "%=", " >>", " << =", "&=", "^ =", and "| =" are not available. This implementation allows very limited structured flow using the "if" and "while" statements. None of the other convenient statements "for," "do while," "switch," "case," "default," and "goto" have been included.

In a small program, many of the above features are not actually necessary, but the lack of ability to logically connect conditions, coupled with a lack of the "for" and "switch-case" operations can make larger programs cumbersome and difficult to program.

As for the I/O functions available, they are rather limited, but usable. Available functions allow single characters and character strings to be read from and written to the user's console. Disk files can be opened and closed, and single byte data can be read in from, or written to any of four open files. Disk I/O is only buffered in 128 byte sectors, however. Another function allows the user to perform CP/M BDOS calls from C.

In addition to the above functions, there is a large amount of generally useful functions inside the compiler source file. It is a shame that these were not pruned from the compiler source and distributed with the other user functions, since they already exist. One such function, called match, compares a specified string with the input stream, and returns true if the two match. This was used in the lexical analysis phase of the compiler, but could just as easily be used in other application programs.

There is a nice feature available that allows inline assembly code to appear in the C source code. Since the compiler generates CP/M 8080 assembler code as its output, inline code can be easily added to the compiler. Any time an "#asm" is encountered in the source, the compiler just takes input and passes it to output with no processing. Once an "#endasm" is encountered, the compiler continues compiling. When I first received the compiler, I had some problems with the assembly code linkages in areas where the manual was unclear, but after a call to The Code Works, I was able to understand it. After my call, they wrote a sample program detailing that linkage, and included it on the distribution disk, so there should be no further problems with other customers.

The Small-C compiler is a one pass compiler (though, technically, the assembler could be considered a second pass). This means that there isn't any checking of defined symbols—everything is just passed on for the assembler to catch. In addition, there is no optimization done on compiled source, either globally or locally. This tends to generate rather inefficient object code, taking more space in memory, and more time to run. In small applications,

PUT YOUR SYSTEM ON THE FOREFRONT OF TECHNOLOGY

PROM BLASTER

PROGRAMS MOST FAMILIES OF EPROMS!
- Accepts 1K/2K/4K or 8K EPROMS
- Extended Device Option
- Phantom Slave Option
- All Programming Characteristics Software Controlled
- Accepts Single or 3 Supply Parts
- Device Address Switch Selectable

Kit Price $199.95

KLUGE CARD

SIMPLIFY YOUR PROJECTS WITH A PROTOTYPE BREADBOARD WITH EXTRAS
- 4 On-Board Pwr. Supplies up to 3 of which can be +5, or ±5, ±12
- Switch Selectable Memory or Device Address
- On-Board Address/Device Decoding
- Bi-Directional Data Bus Buffering
- On-Board Wait States

Bare Board $39.95

6809 SINGLE BOARD COMPUTER

- 2K ROM
- 4K/8K/16K RAM
- 20 Parallel I/O Lines
- RS 232 Interface
- Baud Rates from 110 to 9600
- All Memory/I.O. Relocatable on 4K Boundaries
- 8080 I/O Instructions Memory Mapped (Gives 256 I/O Ports)
- Complete Documentation
- ADSMON Monitor Includes User Callable Functions, Autopatch and more

Kit Price $299.95 (Includes Software)

NOISEMAKER

- Two On-Board Audio Amplifiers for Stereo Sound Effects
- Uses the GIAY 3-8910
- Six Tone Generators
- Two Envelope Generators
- Two Noise Sources
- Four 8-Bit I/O Ports
- Up to Two Wait States

ALL SOUND EFFECTS ARE SOFTWARE CONTROLLED FOR AN ENDLESS VARIETY!
KIT $84.95

MOTHER BOARDS

THREE S-100 MOTHER BOARDS AVAILABLE;
- 3 Slots for Stand-Alone Applications
- 9 Slots for Small Systems
- 18 Slots for Full Size Systems
- Active Termination (9 & 18 Slot Boards Only)
- Extensive Ground Shield
- Breadboard Area
(Write for complete information)

The products shown here are just an example of the quality products and service **Ackerman Digital Systems** offers. **ADS** is an engineering company dedicated to providing the micro computer industry with solid quality products to meet a variety of needs.

All **ADS** products adhere to the **I.E.E.E. 696/S-100 Bus** standard, a bus noted for its widespread acceptance and versatility.

For catalog containing complete hardware and software information, write Ackerman Digital Systems, Inc., 110 N. York Rd., Elmhurst, IL 60126.   TEL 312-530-8992

ads

ACKERMAN
DIGITAL SYSTEMS, INC.

Introduction to C, continued...

this is relatively unimportant, but as the programs increase in size, it can be a real problem.

I found the compiler easy to run, and with no apparent bugs. Unfortunately, compilation is rather sluggish. The run-time package that is included is written in assembler with #asm and #endasm at each end. For a program to assemble, this file must be appended onto the assembler source output of the compiler. The method suggested in the manual is to simply run it through the compiler along with the program being compiled. Doing that is a rather time-consuming process, and I found it much easier to remove the #asm and #endasm statements, and once the compilation was finished, append the run-time software to the end of the compiler output with a text editor. The disadvantage of this is that the entire run-time package is appended, taking up more space in memory.

I was able to compile the compiler source in fourteen minutes on a 4Mhz Z80-A. It took another eight and one half minutes to assemble and 45 seconds to load. After going through all of that, it did run. The source code of the actual compiler was reasonably well written and should be easy to modify. It is also a good example of what a simple, "bare-bones" compiler looks like. In looking at the source, it does not seem that any special parsing or lexical analysis method was used, creating a few additional unnecessary states in code generation.

For anyone looking for a simple compiler to do smaller programming tasks, or for a minimum expense language, the Small-C compiler is a very good choice at $15. For an additional programming tool a few steps above assembler, it is also a good choice. If larger applications are expected, I would not choose this, but go with one of the larger compilers costing from about $145 to $700.

> **Name:** tiny c TWO Compiler
> **Price:** $250
> **Distributed by:**
> tiny c associates
> P.O. Box 269
> Holmdel, NJ 07733

Tiny-C Two Compiler

Tiny-c has been well known for its interpreter of a dialect of the C language, and for its excellent documentation for beginners. After the interpreter was on the market for quite some time, the folks at tiny-c associates felt that although the original tiny-c had quite a following, it was relatively slow for large projects (as an interpreter would be). They decided that a compiler would be very useful.

Now, think about the interpreter compiler combination for a given language. An interpreter is nice and easy to use, and programs can be written, tested, debugged, retested, and so on in the single environment of an interpreter. But interpreters are slow. On the other hand, compilers are faster by at least one order of magnitude, but they are not as simple to work with. You have to enter an editor, type in the program, exit the editor, and attempt to compile. If it won't compile due to syntax errors, you have to reload the editor and edit the program again. When you manage to compile your program, you have to load the linker and debug all of the logic errors. What would offer the best of both worlds would be a

good interpreter to make program development easier, and once it works, a compiler to make the program run faster. For those of you who program in Basic, one combination like this is the Microsoft Basic-80 interpreter, which makes development easier. When you want it to cook right along you can use the Microsoft Basic compiler. However, Basic is not the best language for all things, and it would be convenient to have a somewhat structured language do the same thing. The tiny-c 1 interpreter, combined with tiny-c TWO (the compiler), make a strong step in that direction—although the compiler comes with some very powerful functions not supplied in the interpreter.

Tiny-c is a language that is similar to the C programming language. It is important to realize, however, that tiny-c is a different language from C. It has many constructs and data types to those similar in C, but the syntax and operation are slightly different. Where a program in C is written as:

```
main()
{
      some set of operations
}
```

a tiny-c program is written as:

```
main
[
      some set of operations
]
```

Comments in C are written with a beginning /* and end with a */, such as:

```
getch()
/*
      This routine returns a character from the
      console
*/
{
      some operations
}
```

while comments in tiny-c start with /*, and end with the end of the line, such as:

```
getch
/*
/*      This routine returns a character from the
      console
/*
[
      some operations
]
```

In the C language, statements end with a semicolon ';':

```
range=big-little+1;
```

where in tiny-c, the end of a line or a semicolon signifies the end of a statement, although more than one statement may be on a single line is separated by a semicolon:

```
range = big-little+1
last = last*seed

    -or-

range = big-little+1 ; last = last*seed
```

In C, all compound statements are delimited by the curly braces { and } :

```
{
      x=x-1;
      a=b+c;
      b=b*2-a;
      x=b-a;
}
```

32

MICROSYSTEMS

INTELLIGENT VIDEO I/O FOR S-100 BUS

VIO-X

The VIO-X I/O Interface for the S-100 bus provides features equal to most intelligent terminals both efficiently and economically. It allows the use of standard keyboards and CRT monitors in conjunction with existing hardware and software. It will operate with no additional overhead in S-100 systems regardless of processor or system speed.

Through the use of the Intel 8275 CRT controller with an onboard 8085 processor and 4k memory, the VIO-X interface operates independently of the host system and communicates via two ports. The screen display rate is effectively 80,000 baud.

The VIO-X1 provides an 80 character by 24 line format using a 7 × 9 dot matrix to display the full upper and lower case ASCII alphanumeric 96 printable character set (including true descenders) with special characters for escape and control characters. An optional 2732 character generator is available which allows an alternate 7 × 9 contiguous graphics character set.

The VIO-X2 offers an 80 character by 25 line format using a 9 × 9 dot matrix allowing high-resolution characters to be used. This model also includes expanded firmware for block mode editing.

Both models support a full set of control characters and escape sequences, including controls for video attributes, cursor location and positioning, cursor toggle, light pen location, and scroll speed.

Video attributes provided by the 8275 in the VIO-X include:

- FLASH CHARACTER
- INVERSE CHARACTER
- UNDERLINE CHARACTER or
- ALTERNATE CHARACTER SET
- DIM CHARACTER

The above functions may be toggled together or separately.

The board may be addressed at any port pair in the S-100 host system. Status and data ports may be swapped if necessary. Inputs are provided for parallel keyboard and for light pen as well as an output for audio signalling. The interrupt structure is completely compatible with Digital Research's MP/M

FEATURES

- HIGH SPEED OPERATION
- PORT MAPPED S-100 INTERFACE
- FORWARD/REVERSE SCROLL or
- PROTECTED SCREEN FIELDS
- CONVERSATIONAL or BLOCK MODE
- INTERRUPT OPERATION
- CUSTOM CHARACTER SET
- CONTROL CHARACTERS
- ESCAPE CHARACTER COMMANDS
- INTELLIGENT TERMINAL EMULATION
- TWO PAGE SCREEN MEMORY

VIO-X1 80 × 24 7 × 9 A & T **$295.00**
Conversational & Limited Block Modes
VIO-X2 80 × 25 9 × 9 A & T **$345.00**
Conversational & Block Modes

VIO-X   S-100   I/O INTERFACE

Introduction to C, continued...

In tiny-c they are delimited by square braces [and]:

```
[
    x = x-1
    a = b+c
    b = b*2-a
    x = b-a
]
```

Both allow many levels of nested compound statements.

Functions in C are specified by the function name, a left parenthesis, a list of argument names, and a right parenthesis. This is followed by a type description of each of the arguments:

```
nonprint(c)
char c;
/*
    This function will return true if the character
    is non printable.  It only checks the low order
    seven bits.
*/
{
    return c&128<=32 || c==128;
}
```

It might be called by the statement:

```
if (nonprint(byte))
{
    some statements
}
```

In tiny-c, a function is declared by the name of the function, followed by a type description of each of the arguments.

```
nonprint
char c
/*
/*    This function returns true if the character
/*    is non printable.  It only checks the low order
/*    seven bits.
/*
[
    return c&128<=32
]
```

The function can be called as:

```
if (nonprint(byte))
[
    some statements
]
```

If a function in C returns has no arguments, such as getch, it has to be called like:

```
c=getch();
```

with both parentheses included. In tiny-c, it could be written as:

```
c=getch
```

without the parentheses.

The last major syntactic difference is in the usage of subscripts in arrays. In C, an array is specified by:

```
a=num[index];
```

while in tiny-c:

```
a=num(index)
```

substituting the parenthesis for the square brackets.

For comparison, the following is a small program written in tiny-c and C, taken from the *Tiny-c Newsletter*, Number 1, February, 1981:

```
tiny-c                          C
------                          -----
main                            main()
char a(o)                       char a[o];
[                               {
    int k                           int k;
    while (++k < 10)                while (++k < 10)
    [                               {
        p1""                            p1("");
        pn k                            pn(k);
        ps ""                           ps("");
        putchar a(k)                    putchar (a[k]);
    ]                               }
]                               }
```

The purpose of that newsletter was to explain to users how to convert the tiny-c TWO compiler to a compiler that would compile code closer to standard C. Since the compiler comes with all of the source code, the user only has to make the necessary changes in the code and recompile the compiler. For those who don't wish to make changes themselves, tiny-c associates will be distributing a modified version containing the changes, which they call C-TWO. One of the problems of tiny-c to the experienced C programmer is getting used to and converting old programs to the tiny-c syntax. The C-TWO modification could greatly ease the relearning. For the novice programmer, both to C and to structured languages as a whole, tiny-c is much easier to learn—first with the interpreter, and then with the compiler.

Now that we have looked at the syntactic differences between C and tiny-c, we will look at the functional differences. This is a comparison between the standard C [Kernighan 1978] and tiny-c. Although tiny-c supports many of the features of C, some important items are missing—such as the convenient switch/case statement. In addition, the do/while statement is missing, although the while statement has been implemented. Also missing is the for statement, but this can easily be done in a while loop. A for statement might look like:

```
for( i=1 ; i!=10 ; ++i)
{
    do something
}
```

which can be implemented with a while statement in tiny-c as:

```
i=1
while (i!=10)
[
    do something
    ++i
]
```

It's not as pretty, but it works.

A switch/case statement, on the other hand, is not as easily implemented, nor does it look as nice:

```
switch(num)
{
    case 1:
            statement_1;
            break;          /* exit switch */
    case 2: statement_2;
            break;
    case 3: statement_3;
            break;
}
```

Introduction to C, continued...

This would be done in the form of:

```
if(num==1)
[
    statement_1
]
else
[
    if(num==2)
    [
        statement _2
    [
    else
    [
        if(num==3)
        [
            statment _3
        ]
    ]
]
```

or if not to complex, as:

```
if(num==1) statement_1
else if(num==2) statements_2
else if(num==3) statements_3
```

Tiny-c has the storage classes of static and extern, but does not support auto, register, or typedef. The compiler supports int, long int, and char data types, but does not have floats, doubles, structs, or unions. Tiny-c has a primitive pointer implementation that will address bytes. Normally, in C, when a pointer is declared an int:

```
int *ptr;
```

it points to an integer boundary and, when incremented, it will pass over that integer and point to the next one (usually two bytes later):

```
++ptr;
```

In tiny-c, the pointer will always increment or decrement one byte, regardless of the data type.

Tiny-c also supports many of the unary and binary operators, with the exception of the unary operators type-name, sizeof, and pointer-to (*ptr), in which the latter is implicit when an array is defined. The binary operators "&&" (connective and), " || " (connective or), "=*", "/=", "%=", "+=", "-=", " << =", " >> =", "&=", " ∧ =", and " | =" are also not implemented. The connective and "&&", and or " || " are the most unfortunate losses in that list, not allowing statements such as:

```
if (a==b && c==d) ++c;
```

Tiny-c comes with a wealth of functions in a function library called TCLIB. Most of the relatively common functions listed in standard C are included, plus a set from the tiny-c interpreter known as training functions. These are a set of functions originally set up for beginners, but are a highly convenient set of functions. For instance:

```
pn x            /* Print number x
```

is much more convenient than:

```
fprintf( so, " %d", x)
```

and tend to be very comfortable to use.

Included in the tiny-c TWO package besides the compiler and run-time system is an operating environment called tiny-shell. The tiny-shell is a mised blessing to the compiler. It is loosely based on the UNIX shell developed at Bell

Telephone Laboratories. In order to do any work with the tiny-c compiler, you have to invoke the shell from CP/M by typing "sh" to CP/M's Console Command Processor. This puts you into a new command processor that executes the tiny-c run-time package and user-written tiny-c programs. There are some strong disadvantages to the shell as well as very strong advantages. Using the tiny-c compiler at the present time, you cannot execute programs directly from CP/M—you first must enter the shell. This means that you cannot create turnkey programs for resale without an agreement with tiny-c associates, although in speaking with Tom Gibson, I discover that the cost of a resale license is reasonable. In addition to not being able to execute tiny-c programs from CP/M, you also cannot execute CP/M programs from the tiny-shell. If you chose to do some work typing in the programs with a text editor like Word Star, you have to exit the tiny-shell first by typing a Control-C character. This has the habit of displaying a rather unfriendly message to the terminal:

```
ERROR @ 538 ^C
```

Since the tiny-shell is written in tiny-c, it would have been nice to have the input routine check for a Control-C character and gracefully exit the tiny-shell when encountered. In addition, the operator must become familiar with console input control characters that have slightly different actions in the tiny-shell than in CP/M.

Now for the nice features about the tiny-shell. First, for the programmers who like to define the environment they work in, the tiny-shell (as with everything else in the package) is supplied in source code written in tiny-c. This means that you can modify your environment to your heart's content. In addition, the tiny-shell contains some of the features of I/O redirection found in the UNIX shell. For example, you can write a program that will read in a character from the input device, such as the keyboard, and write it to the output device (the screen). When you run it, it will echo everything typed to the screen.

```
%echo
this is a test
^Z
%
```

Simple, right? Now if you were to use redirection of I/O, and typed:

```
%echo <infile.txt >outfile.txt
%
```

you would copy everything in infile.txt to outfile.txt. The indirection characters "< " for input and " > " for output provide this powerful and helpful feature. For instance, you can compile a program and get a listing of lines and errors displayed on the terminal. You may not want to wait until it gets to line 187 to find out what is wrong, and dig through all of that output. The solution is to redirect the output to a disk file. You then run the output file through a text editor or searching program to find what you want, without the hassle of all that printout. An additional tiny-shell feature allows the user to type multiple commands on a line, hit return, and have it execute a series of commands without intervention.

In addition to the functions built into the tiny shell, later versions of tiny-c TWO come with a set of tiny-c programs that can be compiled into UNIX-like shell functions. These include RM which removes (deletes) a

The tiny-c TWO package is useful for the person who has learned the tiny-c language with the interpreter and wants to be able to do production work, or for introducing the beginner to structured programming.

file, MV which moves (renames) a file, LS which lists the directory and gives a status display, CAT which concatenates ASCII files, HD which gives a hex dump of all or part of a file, DIFF which is a file comparator, ED which is a UNIX-like line oriented text editor, and APRINT which types a file to the terminal.

The tiny-c TWO package comes with the source to everything including the runtime package in assembler, the compiler, the linker, the function library, the tiny-shell, and shell functions. In addition, it comes with a one inch thick manual bound in a bright white three ring binder. The manual describes the tiny-c language in a way that is clear to beginning users. It then goes into a description of how all of the functions are called, how to interface to the machine language, how to use the compiler and tiny-shell, examples of tiny-c programs, internal details of the compiler, and installation to other operating systems. The sections of the manual describing the tiny-c language are very clear, but the section on bringing up the compiler on other machines appears to be written in a totally different manner that may not be clear to the novice. I found the manual to be too verbose for an experienced programmer, requiring the entire book to be scanned to learn how the package works. It would have been nice to have a chapter summarizing the syntax, function calls, and compiler operation.

The tiny-c TWO package is useful for the person who has learned the tiny-c language with the interpreter and wants to be able to do production work, or for introducing the beginner to structured programming. It is also very good for the experienced programmer who likes to modify his system to suit his needs because of the sources for the tiny-shell and compiler. For the experienced C programmer who doesn't wish to be tied to a non-system environment, or who wants the power of such things as switch/case or a for loop statement, I would recommend one of the compilers designed for standard C.

Introduction to C, continued...

Whitesmiths C Compiler

The Whitesmith C compiler was developed by Whitesmiths, Ltd. of New York City. Their compiler was designed to produce executable programs for the 8080/Z80, LSI-11/PDP-11, VAX-11, M68000. They claim that any program written in C from one machine is portable to any of the other machines, provided that no machine-dependent code is written. This compiler is rather large, and by producing an intermediate code, called A-Natural, can be moved from machine to machine. The A-Natural code (which is assembler-like) can then be compiled/assembled into the executable machine code modules for a given machine.

You could say that the Whitesmiths C compiler for CP/M is the Cadillac of the CP/M C compilers. Like a Cadillac, it is very expensive. At $630, it is the most expensive C compiler available to CP/M. And, as a luxury car eats gasoline, this compiler eats memory. If you don't have 60Kbytes, at minimum, don't even bother to attempt to compile the smallest of C programs. Like the Cadillac, it is cumbersome. A compilation and linkage takes much longer with this compiler than with any others. On the other hand, it tends to generate reasonably fast running code. And also like the Cadillac, it is fancy and has all of the flashy options you could want. The Whitesmiths C compiler complies completely with the "standard," with some extra frills added.

The Whitesmiths package for CP/M comes on two single density floppy disks containing the compiler and machine interface code. In addition to the compilation programs, there is a loader, a library manager, a program to examine relocatable modules, and assembler/compiler to translate the A-Natural code to machine code. There is also a program to translate the A-Natural code into Microsoft Macro Assembler code. In addition to utilities, the Whitesmiths C package comes with a library of 87 high-level functions, seven 8080 C callable subroutines and 55 in-line 8080 subroutines. These final routines are used in the operation of the compiler but can also be used by the systems programmer to do some fancy things like multiplying longs by longs. Finally the package includes twelve functions designed for special calls to the CP/M operating system. At a total of 161 functions available, the programmer has truly large library to chose from.

Also included in the Whitesmiths C package is a set of two rather thick manuals entitled "C Compiler Users Manual" and "C Compiler Systems Interface Manual for 8080 Users." The two manuals are interesting to read. They vary from lucid to something rivaling IBM's most obtuse literature. Only after the third or fourth reading does it become possible to understand some of the things available and some of the things that are not.

The Whitesmiths C compiler complies completely with the standard in Kernighan and Ritchie, including floating point, longs, and storage classes. In addition, it includes some things not yet in the standard. These include types unsigned (char short, long) and character constants with more than one character. It also provides command redirection using the CP/M console command processor, allowing such things as:

```
A>read program.prn >other.prn
```

which would take all output of program.prn and place it in other.prn. This feature adds about 4Kbytes to the generated object code. The following program is 6K and generates code for command redirection:

```
main()
{
}
```

while this next one is 2K and has no redirection code:

```
_main()
{
}
```

Finding out about the above feature was a true exercise in documentation reading. Every programmer should try it once (and only once).

The Whitesmiths C compiler provides most of the standard formatted I/O and file functions, but generally under different names. For instance, the common function "printf" is called "putfmt" in this implementation. It also contains functions that operate on arguments of the command line, many string, numerical and data conversion functions.

Although this C compiler produces reasonably fast code, compilation and linking requires a great deal of patience. It compiles in five steps, generating intermediate files, using the CP/M submit facility. The first step is the preprocessor (pp) followed by the parser (p1) and then the 8080 code generator (p2). Once it has generated the A-Natural source code, it must go through the A-Natural assembler/compiler. Finally it must go through a very long linkage process to bring together all the functions and produce executable code.

This compiler has many tradeoffs, but in certain cases it is very valuable to the system developer. On the negative side are documentation that is not entirely clear, the very long time for edit, compile, test, and edit runs, and its requirement for a huge amount of memory. I had to buy a special chip set for my disk controller just to test this compiler. Also the licensing and ordering processes are very comfortable. After signing a very extensive two page license agreement, I received the compiler. From what I understand from the license agreements, even with the payment of over six hundred dollars, Whitesmiths, Ltd. still owns the thing and can require it to be returned any time in the next fifteen years.

On the positive side, the compiler generates fast code which contains the complete C syntax and is thus portable to other machines. This allows a program, such as a compiler, to be developed on one machine (like the PDP-11) and moved either up to the VAX or down to the 8080/Z80 machines (provided you buy three versions of the compiler, of course).

The Whitesmiths C compiler is definitely not the compiler for a casual user. It is however, quite useful and possibly necessary to the person who needs to compile large, production-type programs using the full C syntax. ■

DOS/BIOS Directory And File Conversion In North Star UCSD Pascal (Part II)

by Chris Young

Part II—File Conversion

In this installment we will discuss how to convert the information within the files into standard Pascal data and file types. We will look at some Pascal procedures that facilitate the conversion of North Star data files into similar Pascal data files. This is accomplished by reading the North Star data into standard Pascal types of variables where further processing, including the output of the data to Pascal files, may be done. Three kinds of conversion are discussed: 1) North Star Basic text files into UCSD format text files. 2) North Star Basic string data into Pascal "packed array [..] of char" data. 3) North Star Basic BCD floating point data into Pascal type "real" data.

Text File Conversion

Why should Basic text files be of concern to Pascal users? Prior to the introduction of the new North Word word processing system and Pascal, users did not have a text editing system *specifically* designed for use with North Star systems. True, one can spend over a hundred dollars for CP/M and more for Electric Pencil, Wordstar, or others. However, users of small systems (and users who spent their entire budgets on large systems) may not be able to afford these powerful editors. Even North Word isn't cheap. Many of these users resort to using the Basic program entry editor for their word processing needs. This is accomplished by creating text files full of REM statements. In fact, as long as the user does not try to RUN the files, even the word REM is not necessary. The Basic editor, although limited, can help the user get by on a budget. Most often these files contain documentation of Basic programs, but many other word processing uses can be covered by the Basic editor. A user may have this or many other reasons why he wishes to access Basic text files from Pascal. As his pocketbook recovers, and his editing needs increase, the user may come to the conclusion that for a smaller expenditure for software than is needed for many text editors, he can have the

Chris Young, 3119 Cossell Dr., Indianapolis, IN 46224.

entire UCSD system. This statement is made with the realization that perhaps a significant expenditure for hardware to upgrade to 48K may be necessary. The UCSD system includes a complete Pascal program development system consisting of a compiler, linker, file manager, a Z80 and an 8080 macro assembler, and a powerful screen-oriented text editor. The program described below allows users to access old data files, for whatever reason, for use in this versatile new system.

When a user types a Basic program source text into North Star Basic, the text is not stored exactly as it was entered. Basic recognizes keywords and commands, compacting them into one byte tokens in the range of 128 to 255. The feature not only saves space, but results in faster interpretation of code. When LIST commands are issued, a simple table look-up expands the "crunched" tokens back into their exact original form. However, this space saving is not limited only to Basic statements. Text in string literals and comments (REM statements) are also compacted into one byte tokens wherever possible.

I will now describe the program CVTBAS, found in Listing 4, which performs the function of re-expanding compressed keyword tokens into strings of characters. It creates a UCSD text file or outputs the text to any system device.

First I will examine the procedure GETCHAR which reads the original untyped file, one block at a time, and feeds the characters out one byte at a time. This deblocking routine is the basic tool used in all North Star to UCSD data conversions. GETCHAR comunicates through five variables and an input file which must be declared globally (i.e. at the outermost program level). The file "DOS" is declared as an untyped file (i.e. "var DOS:file;"). LASTCHR is a boolean flag which enables/disables the transmission of characters out of GETCHAR. The flag must be initialized to false at the opening of the input file. After the last byte of the last block has been transmitted, GETCHAR sets LASTCHR to true. Any further calls to GETCHAR returns a null byte. North Star uses a value of one (ASCII character "SOH") as an end of file marker, however,

this is context dependent. GETCHAR cannot set LASTCHR true upon reaching a North Star end of file. The program calling GETCHAR can determine if the byte value of one is part of the data or an end of file mark, and it sets LASTCHR is necessary. The character itself is passed to the calling routines in two global variables: CH and CHINT. CH is of type "char" and CHINT is an "integer" whose value is ORD(CH). BUFR is a "packed array [0..511] of char" which is used as the input buffer. BFPTR is an "integer" which points into BUFR. Each time GETCHAR is called, BFPTR is incremented by one. When BFPTR reaches 512, the buffer is empty and a new block is read. To read the first block, the calling program should initialize BFPTR to 512, forcing a block to be read upon the first call to GETCHAR.

The procedure DOLABEL processes Basic labels which are in the form of line numbers. North Star Basic allows line numbers in the range of 1 through 65535 (64K-1). The line numbers are stored as 16-bit <un-signed> integers. Because UCSD Pascal limits integer variables to -32K to +32K, type "real" variables must be used in the calculations which output the line number. Line numbers appear at the beginning of each line and in certain Basic statements. All in-statement label references are preceded by a token code 154. When a code 154 is received at the beginning of each line, the LABL flag is set to true. The causes the invocation of DOLABEL.

DOLABEL uses five variables. LINENUM is of type "real" and is used to store the line number value. P, also type "real," holds a power of ten to be used to strip off the highest order digits one at a time. LEADZERO is a boolean flag, initialized to true upon entry to DOLABEL, which suppresses the output of leading zeros. LEADZERO is reset to false when the first non-zero digit is processed.

DOLABEL, as well as other routines in the system, requires a character to be in CH and CHINT upon entry to the routine. Likewise, all routines also call GETCHAR upon exit, thereby leaving the next character ready for the next section of processing. Upon entry, DOLABEL initializes LEADZERO and puts the low order byte of the line number in LINENUM. Then GETCHAR is called to get the high order byte. This byte is "shifted left" one byte by multiplying it by 256.0 (remember to use "real" arithmetic). Next add in the low order byte and save it in LINENUM. Because we do not want a decimal point output with the value, we do not use the standard WRITE procedure to output the real valued number. Instead, we strip off the high order digits, one at a time, truncate them to one digit integers, and output them to the output device. The loop runs from 4 down to 0 and the variable P is assigned a value of ten to the fourth down to ten to the zero power. This is used to compute what amounts to integer division and modulo calcuations on a real value. As each digit is stripped off, proceeding left to right, the LEADZERO flag is updated and tested. Non-leading zero digits are output to the output device and to the CONSOLE:. The function of the ISFILE flag will be discussed later.

The main program is initialized by two routines: INIT1 and INIT2. Two routines are needed because there is a limit on the size of procedures. The procedures initialize the table as "KEY:packed array [0..255] of string[10]."

The array elements are all first set to a one character string which is the normal ASCII value for that character. Individual elements are then set to strings of characters which will replace the compacted tokens.

After INIT1 and INIT2 are called, the user is prompted to type the name of the Basic text file to be converted. The response is stored in INPNAME and the file is opened by RESET. The user is then prompted for the output file name. If he hits return, the length of OUTNAME is zero, and ISFILE is set false. If a name is entered, the ISFILE flag is set true, and OUTDEV is created as a text file by REWRITE. The text produced by this program is always sent to the CONSOLE: device. If ISFILE is true, the output is also sent to the file OUTDEV.

All lines begin with a length byte which we will skip over. Next, each line has a label. To process the first one, the LABL flag is set to true. A North Star EOF test is made and we enter a "while not LASTCHR do" loop. Another byte is obtained from GETCHAR because all routines require a character to be in CH and CHINT upon entry. IF LABL is true we call DOLABEL and loop back again. Otherwise, a keyword look-up is done by using CHINT as the index into KEY. The string stored in KEY[CHINT] is output. If CHINT is a carriage return (code 13), then we do an extra GETCHAR to ignore the length byte. This is another place where we test for CHINT=1, meaning a North Star EOF was read. Processing continues as above until LASTCHR is true.

Data File Conversion

Users are much more likely to have a need to convert data files from North Star to Pascal formats. These files may be data bases, tables, mailing lists or any one of a number of other types. The overall format of data files is very application dependent. Rather than trying to give samples of specific data conversions, we will present a general conversion program named CVTDATA (See Listing 5). CVTDATA reads a North Star file and outputs its contents to the CONSOLE: device. The program merely demonstrates the techniques needed. The user must adapt the routines and code sections needed to read the North Star data. Once the data is in Pascal variables, further processing may be done on the data, including output to a new file.

Data in North Star disk files is one of four types: an end of file marker one byte long, a floating point number which occupies five bytes in the standard eight digit version, short strings which occupy 1 to 256 bytes, and long strings which occupy more than 256 bytes of space, with a maximum length which is limited by available memory. CVTDATA reads the input file, determines which of these four types it has encountered, and passes the data to the CONSOLE: with a message telling what type of data was read.

String data is processed by the main program and will be covered later. Floating point data is processed in a separate routine and requires some background information before it may be discussed.

North Star floating point data consists of an eight digit binary coded decimal (BCD) normalized mantissa in four bytes, and an exponent in the fifth byte. The high order bit of the exponent byte contains the sign of the mantissa. The remaining seven bits are an excess 64 exponent base 10. A zero exponent means a value of zero for the entire floating point number. UCSD real variables occupy four bytes. There is a 23-bit binary normalized mantissa with an implied leading 1 bit, a mantissa sign bit, and an 8 bit excess 128 exponent base 2. Zero exponents also denote zero value as in North Star.

The user of North Star/USCD Pascal now has all of the necessary tools to access the contents of DOS/Basic data and text files. With the programs discussed here which allow access to both the directories and the files themselves, the entire package should send the users well on their way to complete conversion to the UCSD system.

What does all of that mean? It means that in North Star, your real variables can have a mantissa plus or minus 9.9999999 and an exponent ten to the plus or minus 64. UCSD has a mantissa plus or minus two to the 24 minus 1, and an exponent two to the plus or minus 128. This works out to an overall range of 9.9999999+ 64 to -9.9999999E-64 for North Star. UCSD has a range of about 0.1701411247E+37 to -0.1701411247E-37, but only about six to seven digits of the mantissa are significant. Note that the UCSD mantissa can exceed 0.17014—if the exponent is of less magnitude. The magnitude of both the mantissa and the exponent of North Star exceeds that of UCSD. This means that both a loss of significance and a possibility of overflow can occur in North Star to UCSD floating point data conversions.

Exponents on the order of ten to the 32 to ten to the 64 may not concern the user unless his data is of a highly scientific nature. However, the loss of significant digits in the mantissa is of concern in circumstances as every day as a seven digit phone number, or dollar amounts over $10,000.00. The user must deal with these problems on a case-by-case basis.

Real data is converted in CVTDATA by a function called CVTREAL. CVTREAL reads North Star real data through our old friend GETCHAR. It returns a variable real parameter containing the value. A boolean overflow flag is the result of the function. Upon entry to CVTREAL the variable VAL is initialized to zero. Four bytes are read via GETCHAR. Each byte is split into two BCD digits giving eight BCD digits total. VAL is multiplied by ten and a digit is added to it for all eight digits. The last digit or two may not add any significance to the mantissa due to the limitation discussed earlier. Because of the way the mantissa is built up, it will be necessary to divide by 1.0E08 later in the processing. The fifth byte is

read and the mantissa sign bit is striped. The excess 64 is subtracted from the exponent, and its sign is determined. Overflow conditions are tested, and if they occur the value of VAL is set to a maximum. The exponent is either multiplied or divided into VAL and VAL is normalized. As usual, GETCHAR is called again to make ready for the next routine.

The main program begins by prompting the user for the input file name. Recall that the program is only a model which outputs to the CONSOLE: device, so no output name is requested. LASTCHR and BFPTR are initialized. GETCHAR obtains the first character.

While there are still characters to process, the following actions occur. If CHINT is greater than 15 real data is to be processed. A messages that a number is being processed is printed, and an overflow indication (if necessary) follows that. Finally, the value obtained by CVTREAL is printed. If the original value of CHINT was less than or equal to 15, then a "case" statement is used to process the other options. Option one is an end of file. When North Star EOF is read, a message is printed and LASTCHR is set to true. This causes an exit from the "while" loop and the program terminates. Case two is a long string. Long strings have a two byte, low order byte first, length field. GETCHAR gets the first byte and it is stored in COUNT. GETCHAR gets the second byte, which is multiplied by 256 and added to COUNT. The third case is a short string. The length filed is a one byte field which is read by GETCHAR and saved in COUNT. After leaving the "case" statement, the string is output

by a loop whose index runs from 1 to COUNT. The string is output to the CONSOLE: one byte at a time. CHINT values of zero or ten (i.e. nulls and line feeds) are skipped. This is not a necessary feature and can be eliminated if desired. The "one character at a time" output can be replaced by code which puts the characters in a packed array or string for further processing if needed.

The principals demonstrated in CVTDATA can be applied to any North Star data where the precision restrictions are not of concern. Listing 6 is a procedure called WRITREAL which can be incorporated into CVTDATA to replace CVTREAL with slight modification to CVTDATA. WRITREAL reads an eight digit North Star value and outputs it to the CONSOLE: in standard scientific notation. The user may be able to modify this routine to suit the need to extract all of the significance of real data. How he manages to further process this data in Pascal is up to him.

Summary

Some of the many advanced features now available to North Star Pascal users have been demonstrated in the directory conversion programs. The user of North Star/UCSD Pascal now has all of the necessary tools to access the contents of his DOS/Basic data and text files. With the programs discussed above which allow access to both the directories and the files themselves, the entire package should send the users well on their way to complete conversion to the UCSD system. So go to it! ∎

Listing starts on page 44.

================ Listing 4 ================

```
            (* Program to produce a U.C.S.D. text  *)
            (*  file from a North Star BASIC text.*)
            (*  Written Aug. '80                   *)
            (*      by Chris Young                  *)
            (*         3119 Cossell Drive           *)
            (*         Indianapolis IN 46224        *)
            (*         (317)-291-5376               *)
program CVTBAS;
  var KEY            (* used to map keywords into strings *)
          :packed array[0..255]of string[10];
      CH:char;
      BUFR:packed array[0..511] of char;
      I, J,          (* indicies *)
      CHINT,         (* contains ORD(CH) *)
      BFPTR          (* pointer into input buffer *)
          :integer;
      LASTCHR,       (* is true after last character is processed *)
      LABL,          (* if true, next 2 bytes contain a line number *)
      ISFILE         (* output flag *)
          :boolean;
      OUTNAME,       (* name of output device or file, if null then *)
                     (*  output to console only                    *)
      INPNAME        (* name of BASIC file to be read *)
          :string[20];
      BASIC          (* file of BASIC text to be converted to *)
                     (*  U. C. S. D. text                    *)
          :file;
      OUTDEV         (* file for converted text *)
          :text;
  procedure INIT1;
   begin
    (* initialize array of keyword tokens to default to   *)
    (* their ASCII values                                 *)
    for I:=0 to 255 do
     begin
      KEY[I]:=' ';              (* make strings of length 1 *)
      KEY[I][1]:=CHR(I); (* put in the ASCII value   *)
     end;
    (* now fill in keywords which are not one-to-one with ASCII *)
    KEY[128]:='LET';    KEY[129]:='FOR';
    KEY[130]:='PRINT';  KEY[131]:='NEXT';
    KEY[132]:='IF';     KEY[133]:='READ';
    KEY[134]:='INPUT';  KEY[135]:='DATA';
    KEY[136]:='GOTO';   KEY[137]:='GOSUB';
    KEY[138]:='RETURN'; KEY[139]:='DIM';
    KEY[140]:='STOP';   KEY[141]:='END';
    KEY[142]:='RESTORE'; KEY[143]:='REM';
    KEY[144]:='FN';     KEY[145]:='DEF';
    KEY[146]:='!';      KEY[147]:='ON';
    KEY[148]:='OUT';    KEY[149]:='FILE';
    KEY[150]:='EXIT';   KEY[151]:='OPEN';
    KEY[152]:='CLOSE';  KEY[153]:='WRITE';
    (* code 154 means a line number label follows *)
    KEY[154]:='';       KEY[155]:='CHAIN';
    KEY[156]:='LINE';   KEY[157]:='DESTROY';
    KEY[158]:='CREATE'; KEY[159]:='ERRSET';
    KEY[160]:='RUN';
   end; (* INIT2 *)
  (* the INIT procedure is too large so split it in two *)
  procedure INIT2;
   begin
    KEY[161]:='LIST';   KEY[162]:='MEMSET';
    KEY[163]:='SCR';    KEY[164]:='AUTO';
    KEY[165]:='LOAD';   KEY[166]:='CONT';
    KEY[167]:='APPEND'; KEY[168]:='REN';
    KEY[169]:='NSAVE';  KEY[170]:='SAVE';
    KEY[171]:='BYE';    KEY[172]:='EDIT';
    KEY[173]:='DEL';    KEY[174]:='PSIZE';
    KEY[175]:='CAT';    KEY[176]:='STEP';
    KEY[177]:='TO';     KEY[178]:='THEN';
    KEY[179]:='TAB';    KEY[180]:='ELSE';
    KEY[181]:='CHR$';   KEY[182]:='ASC';
    KEY[183]:='VAL';    KEY[184]:='STR$';
    KEY[185]:='NOFNDMARK';      KEY[186]:='INCHAR$';
    KEY[187]:='FILE';   KEY[224]:='(';
    KEY[255]:=';';      KEY[226]:='*';
    KEY[227]:='+';      KEY[228]:='[';
    KEY[229]:='-';      KEY[231]:='/';
    KEY[236]:='AND';    KEY[237]:='OR';
    KEY[239]:='>=';     KEY[240]:='<=';
    KEY[241]:='<>';     KEY[244]:='<';
    KEY[245]:='=';      KEY[246]:='>';
    KEY[247]:='NOT';    KEY[198]:='INT';
    KEY[204]:='LEN';    KEY[205]:='CALL';
    KEY[206]:='RND';    KEY[202]:='SGN';
    KEY[203]:='SIN';    KEY[210]:='ATN';
    KEY[216]:='FREE';   KEY[217]:='INP';
    KEY[218]:='EXAM';   KEY[219]:='ABS';
    KEY[220]:='COS';    KEY[221]:='LOG';
    KEY[222]:='EXP';    KEY[223]:='TYP';
   end; (* INIT2 *)
  procedure GETCHAR;
   (* This procedure gets a character from the file BASIC and    *)
   (* returns the character in the global variables CH and CHINT.*)
   var N:integer;
   begin(*GETCHAR*)
     if not LASTCHR              (* if last character has *)
                                 (* been read, don't try again *)
      then begin
           if BFPTR=512  (* buffer is empty, get new buffer *)
             then if EOF(BASIC)
                  then begin       (* no more buffers in file *)
                        BFPTR:=0;
                        BUFR[0]:=CHR(0);
                        LASTCHR:=true; (* this shuts off GETCHAR *)
                                       (* even if N* EOF has not *)
                                       (* been read             *)
                       end (* if EOF(BASIC) *)
                  else begin
                        N:=BLOCKREAD(BASIC,BUFR,1);
                        BFPTR:=0;
                       end; (* if not EOF(BASIC) *)
           CH:=BUFR[BFPTR]; CHINT:=ORD(CH);
           BFPTR:=BFPTR+1;
          end; (* if not LASTCHR *)
   end; (* GETCHAR *)
  procedure DOLABEL;
   (* This procedure is called after every occurence of the      *)
   (*  label flag which is code 154, and at the begining of every *)
   (*  line. It reads 2 bytes and outputs the integer value as a  *)
   (*  string of ASCII characters.  The value is interpreted as a *)
   (*  16 bit unsigned integer from 0 to 65535.  Note this is     *)
   (*  outside the range of Pascal integer variables, so real     *)
   (*  variables are used.                                        *)
   var LINENUM,  (* the value of the label as a real number       *)
       P         (* real power of ten   *)
          :REAL;
       LEADZERO  (* supresses output of leading zeroes         *)
          :boolean;
   begin
    LEADZERO:=true;
```

```
          LINENUM:=CHINT;        (* low order byte  *)
          GETCHAR;               (* high order byte *)
          LINENUM:=LINENUM + 256.0 * CHINT;
          LABL:=false;           (* am no longer reading label *)
          for I:=4 downto 0 do
            begin
            P:=PWROFTEN(I);
            J:=TRUNC(LINENUM/P);        (* divide out high order digit *)
            LINENUM:=LINENUM - J * P;
            LEADZERO:=LEADZERO and (J=0);   (* LEADZREO true until J<>0 *)
            if not LEADZERO then
              begin
              if ISFILE then WRITE(OUTDEV,J);
              WRITE(J);
              end; (* if not LEADZERO *)
            end; (* for I *)
          end; (* DOLABEL *)
begin (* MAIN *)
    INIT1;    INIT2;
    WRITE('Type BASIC File-name:'); READLN(INPNAME);
    RESET(BASIC,INPNAME);
    WRITE('Type new file name:');    READLN(OUTNAME);
    if LENGTH(OUTNAME)>0 then begin
                              REWRITE(OUTDEV,OUTNAME);
                              ISFILE:=true
                              end
                         else ISFILE:=false;
    LABL:=true;    (* file starts with a label *)
    LASTCHR:=false;
    BFPTR:=512;    (* this means buffer is empty *)
    GETCHAR;       (* skip length byte *)
    LASTCHR:=LASTCHR or (CHINT=1);     (* test for N* EOF *)
    while not LASTCHR do
      begin
      GETCHAR;
      if LABL
        then DOLABEL
        else begin
             LABL:=(CHINT=154) or (CHINT=13);
             if ISFILE then WRITE(OUTDEV,KEY[CHINT]);
             WRITE(KEY[CHINT]);
             if CHINT=13 then
               begin
               GETCHAR;        (* skip length byte *)
               LASTCHR:=LASTCHR or (CHINT=1);   (* test for N* EOF *)
               end;(*if CHINT=13*)
             end;(*if LABL*)
      end;(*while*)
    CLOSE(OUTDEV,LOCK);
    end.
```

═══════════ Listing 5 ═══════════

```
          (* Program to demonstrate North Star  *)
          (*   to U.C.S.D. Pascal data conv.     *)
          (*   Written Aug. '80                  *)
          (*      by Chris Young                 *)
          (*         3119 Cossell Drive          *)
          (*         Indianapolis IN 46224       *)
          (*              (317)-291-5376         *)
program CVTDATA;
    var CH            (* character returned by GETCHAR *)
```

```
                     :char;
        BUFR        (* input buffer used by GETCHAR    *)
                    :packed array[0..511] of char;
        I, J,       (* indicies *)
        CHINT,      (* contains ORD(CH) returned by GETCHAR *)
        COUNT,      (* length of string *)
        BFPTR       (* pointer into input buffer *)
                    :integer;
        LASTCHR     (* is true after last character is processed *)
                    :boolean;
        VALU        (* temp real *)
                    :real;
        DOS         (* N* DOS format data file for input *)
                    :file;
        INPNAME     (* name of DOS file to be read *)
                    :string[20];
procedure GETCHAR;
    (* This procedure gets a character from the file DOS and *)
    (*   returns the character in the global variables *)
    (*   CH and CHINT.               *)
    var N:integer;
    begin(* GETCHAR *)
      if not LASTCHR          (* if last character has *)
                              (* been read, don't try again *)
        then begin
             if BFPTR=512   (* buffer is empty, get new buffer *)
               then if EOF(DOS)
                    then begin       (* no more buffers in file *)
                         BFPTR:=0;
                         BUFR[0]:=CHR(0);
                         LASTCHR:=true; (* this shuts off GETCHAR *)
                                        (* even if N* EOF has not *)
                                        (* been read              *)
                         end (* if EOF(DOS) *)
                    else begin
                         N:=BLOCKREAD(DOS,BUFR,1);
                         BFPTR:=0;
                         end; (* if not EOF(DOS) *)
             CH:=BUFR[BFPTR]; CHINT:=ORD(CH);
             BFPTR:=BFPTR+1;
             end; (* if not LASTCHR *)
      end; (* GETCHAR *)
function CVTREAL(var VAL:real): boolean;
    var EXPSGN,   (* sign of exponent, +1 or -1      *)
        HI,       (* high order 4 bit BCD value of a byte *)
        LO,       (* low order 4 bit BCD value of a byte *)
        I,        (* index *)
        EXPON     (* exponent as integer *)
                  :integer;
        A         (* real value of ten to the EXPON *)
                  :real;
        OVFL:boolean;
    begin
    VAL:=0;
    for I:=0 to 3 do     (* read 4 bytes, 8 BCD digits *)
      begin
      HI:=CHINT div 16; (* high order 4 bits *)
      LO:=CHINT mod 16; (* low order 4 bits *)
      VAL:=VAL * 100 + HI * 10 + LO;     (* build up value *)
      GETCHAR;
      end;
    if CHINT>127 then    (* high order bit of 5th byte *)
                         (* is sign of mantissa        *)
```

```
    begin
      CHINT:=CHINT-128;
      VAL:=-VAL;
    end;
 EXPON:=CHINT;
 if EXPON<>0              (* zero exponent means zero value *)
   then
     begin
       EXPON:=EXPON-64; (* exponent is excess 64 *)
       if EXPON>0 then EXPSGN:=1 else EXPSGN:=-1;
       EXPON:=ABS(EXPON);
       (* Note:  0.1701411247E+-37 is the limit before overflow *)
       OVFL:=(EXPON>37)or((EXPON=37)and(VAL>17014112.0));
       if OVFL then begin EXPON:=37; VAL:=17014112.0 end;
       A:=PWROFTEN(EXPON);
       if EXPSGN<0 then A:=1.0/A;
       (* VAL is on the order of 1E8 so normalize it and *)
       (*   multiply in exponent.                         *)
       VAL:=VAL/1.0E8 * A;
     end (* if EXPON<>0 *)
   else
     begin
       VAL:=0; OVFL:=false;
     end; (*if EXPON=0*)
 GETCHAR;      (* always leave a character in CH and CHINT *)
 CVTREAL:=OVFL;
end;(*CVTREAL*)
begin (* CVTDATA *)
 WRITE('Type DOS data file name:'); READLN(INPNAME);
 RESET(DOS,INPNAME);
 LASTCHR:=false;
 BFPTR:=512;     (* this means buffer is empty *)
 GETCHAR;        (* skip length byte *)
 LASTCHR:=LASTCHR or (CHINT=1);     (* test for N* EOF *)
 while not LASTCHR do
  begin
    if CHINT>15 (* this means it is a real value *)
      then begin
          WRITE('Number =');
          if CVTREAL(VALU) then WRITE('**** OVERFLOW **** ');
          WRITELN(VALU);
        end (* if CHINT>15 *)
    else
       begin
         case CHINT of
           1: begin
               WRITELN('EOF');
               LASTCHR:=true;
             end;
           2: begin
               WRITE('Long string ');
               GETCHAR;
               COUNT:=CHINT;
               GETCHAR;
               COUNT:= 256 * COUNT + CHINT;
               GETCHAR;     (* always leave a character *)
                            (*   in CH and CHINT        *)
             end;
           3:begin
               WRITE('Short string ');
               GETCHAR;
               COUNT:=CHINT;
               GETCHAR;     (* always leave a character *)
                            (*   in CH and CHINT        *)
             end;
```

```
          end; (* case *)
         if not LASTCHR then
           begin
             WRITE('"');
             for I:=1 to COUNT do
               begin
                 if (CHINT<>0) and (CHINT<>10) then WRITE(CH);
                 GETCHAR;
               end; (*for I:=1 to CHINT *)
             WRITELN('"');
           end; (* if not LASTCHR *)
       end; (* if CHINT>15 then ... else *)
  end; (* while not LASTCHR *)
end. (* CVTDATA *)
```

===**Listing 6**===

```
              (* Program to output North Star 8 dig.*)
              (*   real to CONSOLE:                 *)
              (*   Written Aug. '80                 *)
              (*       by Chris Young               *)
              (*        3119 Cossell Drive          *)
              (*        Indianapolis IN 46224        *)
              (*        (317)-291-5376              *)
procedure WRITREAL;
(* Reads North Star 8 digit real and outputs it to CONSOLE:   *)
  var I,        (* index *)
      EXPON     (* exponent as integer *)
          :integer;
      MANTSGN   (* if true, mantissa is negative *)
          :boolean;
      DIGITS    (* array of 8 digits *)
          :packed array[0..7] of char;
begin
  for I:=0 to 3 do     (* read 4 bytes, 8 BCD digits *)
   begin
     DIGITS[I*2]:=CHR((CHINT div 16)+48);    (* high order 4 bits *)
     DIGITS[I*2+1]:=CHR((CHINT mod 16)+48); (*  low order 4 bits *)
     GETCHAR;
   end;
  if CHINT>127  (* high order bit of 5th byte is sign of mantissa *)
    then
     begin
       CHINT:=CHINT-128;
       MANTSGN:=true;
     end
   else MANTSGN:=false;
  EXPON:=CHINT;
  if EXPON<>0            (* zero exponent means zero value *)
    then
     begin
       EXPON:=EXPON-64; (* exponent is excess 64 *)
       if MANTSGN then WRITE('-') else WRITE('+');
       WRITE('0.');       (* output first digit and decimal pt. *)
       for I:=0 to 7 do WRITE(DIGITS[I]);   (* output digits *)
       WRITE('E',EXPON);
     end (* if EXPON<>0 *)
    else
     WRITE('0.000000E+00');
  GETCHAR;    (* always leave a character in CH and CHINT *)
end; (*WRITREAL*)
```

Virtual Segment Procedures under UCSD Pascal

by Jon Bondy

One of the nice features of UCSD Pascal is its support of segment procedures. A segment procedure is like any other Pascal procedure except that whenever it is called (except for recursive calls), it is loaded from disk memory prior to being executed; and after it exits, the memory is reclaimed. In fact, the segment procedure is loaded onto the stack since the pattern of memory use is nested in a very stack-like manner.

> *A segment procedure is like any other Pascal procedure except that whenever it is called (except for recursive calls), it is loaded from disk into memory prior to being executed; and after it exits, the memory is reclaimed.*

Segment procedures allow the programmer to manage memory resources explicitly and conveniently, and really are a form of overlay. In large programs, it is not uncommon to dedicate a segment procedure to initialization, since that code need not reside in memory after the program starts. Segment procedures may have internal procedures and functions, all of which are loaded with the segment procedure code, allowing functional groups of routines

Jon Bondy, Box 148, Ardmore, PA 19003.

> *Segment procedures allow the programmer to manage memory resources explicitly and conveniently, and really are a form of overlay.*

to be brought into memory in a single operation, and executed.

Unfortunately, UCSD Pascal allows the programmer access to only six segment procedures under normal circumstances. This is fine for small programs (under 3000 lines), but when one starts to get serious about an application, more segment procedures are necessary. One reason for this is that the UCSD Pascal separate compilation construct (UNITs) uses one of these six segment procedure slots even if it is not loaded into memory dynamically, wasting this scarce resource. This situation should improve very soon (perhaps by the time this article is printed) because Softech is planning to announce features in UCSD Pascal Version IV.O to solve some of these problems.

I work for a company (Energy Data Systems) which has been trying to do some fairly complex applications in UCSD Pascal, and we ran into the "segment barrier" in the spring of 1980. We considered modifying the UCSD operating system since we had source for it, but decided to try to stick to solutions which required as few

48   MICROSYSTEMS

modifications to operating system code as possible. After some thought, I came up with an interim means of ameliorating the problem, which I will describe in this article.

When a segment procedure is called by the UCSD Pascal p-machine, a special op-code is used. This op-code first looks in an operating system segment table to see if the indicated procedure has been called more than zero times (reference count, in case it is a recursive call and the code need not be reloaded). If it is already in memory (reference count > 0), it is executed like any other procedure; if not, the op-code looks in the segment table for the block number on the disk where it can find the code for the segment procedure, and the number of bytes in the code. It then loads the code onto the stack and calls the procedure, incrementing the reference count to one. Upon exit, the segment procedure return op-code decrements the reference count by one, and clears the stack back up to where it was prior to the call if the count has become zero.

My idea was to somehow write data into the segment tables prior to each segment procedure call in such a way that when the above op-code was invoked, it would find data describing different segment procedures each time; it would in fact be faked into loading different code segments into memory for each call, even though the same segment procedure was being called each time.

In order to do this, I first had to find where the segment tables were located in memory. The first global variable declared in the UCSD Pascal operating system code is a pointer to a special record, called the system communication record, or SYSCOMREC. The segment tables are a part of this record. There is a special switch in the UCSD Pascal compiler (the 'U' switch) which, when its value is '-', causes programs to function very differently than usual. The main program does not execute at all, but rather the first segment procedure declared in the program executes instead. Also, the variable definitions are "aliased" on top of the definitions for the operating system (like an EQUIVALENCE in Fortran), allowing the program to read and write those variables. Usually, one uses exactly the same variable definitions as are used when the operating system is compiled, in order that the program agree with operating system definitions. In this case, however, I simply made my own definitions, since all I wanted was to determine the value of the pointer to SYSCOMREC. Since 1.5 stores pointers as actual memory addresses, by printing the value of the pointer I could determine where in memory SYSCOMREC was stored.

I wrote the following program and was able to locate SYSCOMREC. It was part trial and error, since after I thought I had found SYSCOMREC the first time, I was forced to look through dumps to figure out why I had been wrong. Anyway, the program below will tell you

```
($U-)
program find;
  var
    i : integer; ( aliased to tsyscomrec )

  segment procedure findsyscom;
    begin
      writeln('Syscomrec is located at address ',i,' decimal.');
    end;

begin
end.
```

where in memory any SYSCOMREC is under Version I.5 (it is at location 718 [decimal] for my Z-80 version of I.5), and probably under Version II.0. I have not tried it with III.0.

I then needed to figure out how far from the start of SYSCOMREC the segment tables started. Fortunately, the UCSD Pascal operating system variable definitions, found in a file called GLOBALS.TEXT, were distributed with UCSD Pascal Versions 1.4 and 1.5, so I had them at my disposal. [The GLOBALS.TEXT file, along with all other UCSD Pascal source code, is a copyright of the University of California at San Diego; some of the GLOBALS.TEXT file is presented in this article (the portions of Pascal source in capital letters) with the permission of Softech Microsystems Inc., their licensee.] I am informed that the SYSCOMREC data layout has not changed with the various UCSD Pascal Versions, so you should be able to locate the segment tables at 96 bytes past the start of your SYSCOMREC. This means that in writing to location 814 in my memory, one would be writing on the first byte of the segment tables.

The segment tables are defined as follows:

```
type

  segdesc = record
    diskaddr : integer; ( absolute block number on disk )
    codeleng : integer; ( in bytes )
    end;

  segtab = array [segrange] of record
    unit : unitnum; ( disk unit number (an integer) )
    codedesc : segdesc; ( as above )
    end;
```

where "segrange" is the number of segment procedures defined for the UCSD system being used ("0..15" in the case of Version 1.5). If one declared a pointer "segptr" which pointed to a record of type "segtab", one could refer to that record as "segptr ↑ ", to the unit (disk drive) on which the code for the third segment procedure was located as "segptr ↑ [3].unit", and to the number of bytes in the code for that procedure as "segptr ↑ [3].codedesc. codeleng". One could establish the correct value in that pointer with the following record definition:

```
var

  alias : record case boolean of
    true : (i : integer);
    false: (p : tsegtab);
    end;
```

This record definition states that one will either use the storage for the record alias as an integer (denoted "alias.i") or as a pointer to a variable of type "segtab" (denoted "alias.p"). Since the same storage is used for both values, if one were to write into the integer part, one could then use that value as a pointer; one could trick Pascal into thinking that a variable of type "segtype" was being pointed to.

By stating "alias.i := 718 + 96", I could access the actual operating system segment table as "alias.p ↑". I could then write into the segment table entries of any segment which I chose, forcing the system to load the code I wanted to when I called the appropriate procedure.

Suppose, for instance, that I knew that I wanted to

execute each of a series of ten segment procedures which were located on "unit[i]", had length "length[i]" bytes, and started at disk block number "block[i]". I could write a procedure to perform a call to the "i-th" such segment procedure as follows:

```
Program test;

   type
      segdesc = record
         diskaddr : integer; ( absolute block number on disk )
         codelens : integer; ( in bytes )
         end;
      segtab = array [segrange] of record
         unit : unitnum; ( disk unit number )
         codedesc : segdesc; ( as above )
         end;

   var
      alias : record case boolean of
         true : (i : integer);
         false: (p : ↑segtab);
         end;
      unit : array[1..10] of integer;
      length : array[1..10] of integer;
      block : array[1..10] of integer;
      i : integer;

   segment Procedure virtual;
      besin ( need not have any code, since it will never execute --
             the other ten segment procedures will execute instead )
      end; ( virtual )

   procedure dovirtual(i : integer);
      besin
      ( set up to call virtual segment )
      alias.p↑[10].codeunit := unit[i];
      alias.p↑[10].codedesc.codelens := lensth[i];
      alias.p↑[10].codedesc.diskaddr := block[i];
      ( call virtual segment loaded above )
      virtual;
      end; ( dovirtual )

besin
( set up pointer to real segment table )
alias.i := 718 + 96;
( call the ten virtual procedures )
for i := 1 to 10 do dovirtual(i);
end. ( test )
```

The above program will work, but has a few problems which make it a bit awkward to use. First, procedure calls which previously looked like a nice name now are reduced to a cryptic statement like "dovirtual(3)". This problem can be corrected by creating constants at the start of the program with values from 1 to 10, and calling "dovirtual" with those constant values; a procedure to clear the screen might then be called as "dovirtual (clrscreen)", a significant improvement.

The other problem is that it is *not easy* to find out some of the information which I so casually stated would be found in the length and block arrays. To do so involves reading the segment table of the code file in which one of the virtual segments exists. Unfortunately, the disk address information stored in the segment tables of a code file is slightly different than that stored in the operating system's segment tables. The normal code file disk addresses are relative to the start of the entire code file; the system addresses are the absolute disk block number. This means that in order to convert the data in the code file's segment table into useful information, we must also know the absolute address of the start of the code file. And in order to determine this, we must read (and understand) the directory of the disk. Whew!...

Taking it a step at a time, the format of a UCSD Pascal disk directory is given below. It is a portion of the UCSD GLOBALS.TEXT file mentioned earlier (and is copyrighted by UCSD).

```
CONST
   MAXUNIT = 12;        (*MAXIMUM PHYSICAL UNIT # FOR UREAD*)
   MAXDIR = 77;         (*MAX NUMBER OF ENTRIES IN A DIRECTORY*)
   VIDLENG = 7;         (*NUMBER OF CHARS IN A VOLUME ID*)
   TIDLENG = 15;        (*NUMBER OF CHARS IN TITLE ID*)
   FBLKSIZE = 512;      (*STANDARD DISK BLOCK LENGTH*)
   DIRBLK = 2;          (*DISK ADDR OF DIRECTORY*)
   MAXSEG = 15;         (*MAX CODE SEGMENT NUMBER*)

TYPE

   DATEREC = PACKED RECORD
      MONTH: 0..12;              (*0 IMPLIES DATE NOT MEANINGFUL*)
      DAY: 0..31;                (*DAY OF MONTH*)
      YEAR: 0..100               (*100 IS TEMP DISK FLAG*)
      END (*DATEREC*) ;

   UNITNUM = 0..MAXUNIT;
   VID = STRING[VIDLENG];

   DIRRANGE = 0..MAXDIR;
   TID = STRING[TIDLENG];

   FILEKIND = (UNTYPEDFILE,XDSKFILE,CODEFILE,TEXTFILE,
      INFOFILE,DATAFILE,GRAFFILE,FOTOFILE,SECUREDIR );

   DIRENTRY = RECORD
      DFIRSTBLK: INTEGER;              (*FIRST PHYSICAL DISK ADDR*)
      DLASTBLK: INTEGER;              (*POINTS AT BLOCK FOLLOWING*)
      CASE DFKIND: FILEKIND OF
         SECUREDIR,
         UNTYPEDFILE:
            (DVID: VID;                (*ONLY IN DIR[0]...VOLUME INFO*)
                                       (*NAME OF DISK VOLUME*)
             DEOVBLK: INTEGER;         (*LASTBLK OF VOLUME*)
             DNUMFILES: DIRRANGE;      (*NUM FILES IN DIR*)
             DLOADTIME: INTEGER;       (*TIME OF LAST ACCESS*)
             DLASTBOOT: DATEREC);      (*MOST RECENT DATE SETTING*)
         XDSKFILE,CODEFILE,TEXTFILE,INFOFILE,
         DATAFILE,GRAFFILE,FOTOFILE:
            (DTID: TID;                (*TITLE OF FILE*)
             DLASTBYTE: 1..FBLKSIZE;   (*NUM BYTES IN LAST BLOCK*)
             DACCESS: DATEREC)         (*LAST MODIFICATION DATE*)
      END (*DIRENTRY*) ;

   var
      directory : array[dirrange] of direntry;
```

Given the above definitions, after a bit of studying we can see that each directory entry contains the disk address (in blocks) of the first and last blocks of the file for which it is an entry. If the entry is a volume ID entry, it also contains the volume (diskette) name, the number of blocks on the volume (deovblk), the number of files on the volume (dnumfiles) and the time (date) when it was last accessed. Normally one finds this entry as the first entry in the directory (or "directory[0]"). If the entry is a normal file entry, it contains the file ID (name), the number of bytes in the last block which really contain data, and the date of the last modification to the file. The above definitions also tell us where to find the directory, namely at "dirblk", or block two.

We can read the directory into memory with the statement

```
unitread( unitnum,directory,sizeof(directory),dirblk );
```

This statement uses two UCSD intrinsics, the "sizeof" function and the "unitread" procedure. The former returns the number of bytes in a given data structure; in our case it is in the number of bytes in the directory. The unitread procedure reads from unit (disk drive) "unitnum" into memory at the address of the directory structure for "sizeof(directory)" bytes starting at absolute disk block dirblk.

With this information in memory, we can search the directory for an entry of a given file. Suppose that we wanted to see if the file "sample.code" were in the directory. We could say

52

Virtual Segment Procedures, continued...

```
numfiles := directory[0].dnumfiles; ( number of valid entries )
i := 0; ( will index into directory )
done := false; ( will become 'true' when we are done )
while (i < numfiles) and not done do begin
   i := i + 1;
   if (directory[i].dtid = 'sample.code') then done := true;
   end;
```

If we come out of this loop with "done" having a value of "true," then we have indeed found an entry for the required file; if not, the file was not present. If it was found, then the absolute address of the first block of the file on disk is simply "directory[i].dfirstblk". We have then found the first thing we needed, the absolute disk address of the code file. We now need to get the code length and relative disk address information from the code file's segment tables.

The format of the first block of a code file is as shown below, where the previous type definitions hold as before.

```
sestable : array[segrange] of segdesc;
```

If we wanted to find the necessary information about a particular code file, we could use the UCSD Pascal system intrinsic procedure unitread to read the segment table data from that file into the above data structure. Then, if we wanted the information about the tenth segment procedure in a code file, we could simply use the tenth element of that array. For example, to determine the length and block data for segment procedure number ten in a file called "sample.code" on unit five, one could do the following:

```
program find;
   var
      codefile : file;
      sestable : array[segrange] of segdesc; ( 'segdesc' defined as before )
      i : integer;
begin
   unit := 5;
   ( read in segment table -- assume 'directory' read in as above first,
     and 'directory[i]' is data for file which is of interest to us )
   unitread(unit, sestable, sizeof(sestable), director[i].dfirstblk);
   length := sestable[10].codelen;
   block := sestable[10].diskaddr; ( relative block number )
   block := block + directory[i].dfirstblk; ( absolute block number )
   end. ( find )
```

So, we finally have the entire ball of wax. We can read a UCSD disk directory to find out where on disk a file is stored; we can read and understand the information in the segment table of a code file; we can convert the relative block numbers in the code file segment tables to absolute block numbers; and we can use that information to invoke virtual segment procedures. To tie everything together, you will find below a single program which does it all. The procedure "virtinit" initializes an internal table of unit, length, and block information before the program really gets going. It does this by reading in the disk directory and calling "virtlink" for each virtual segment procedure which it will call later. The virtlink procedure looks in the directory for the code file requested and puts the required information in the internal table. The main program then invokes the virtual segments as a test. Following this program is a sample of one of the programs which defines a virtual segment procedure. Note that in both the program which calls the virtual segments and the program which defines one of them, the virtual segments are defined as the first executable code in the file. This is necessary, since the technique which I have shown requires that both segment procedures have the same segment numbers. ∎

```
program vsestest;

CONST
   MAXUNIT = 12;         (*MAXIMUM PHYSICAL UNIT # FOR UREAD*)
   MAXDIR = 77;          (*MAX NUMBER OF ENTRIES IN A DIRECTORY*)
   VIDLENG = 7;          (*NUMBER OF CHARS IN A VOLUME ID*)
   TIDLENG = 15;         (*NUMBER OF CHARS IN TITLE ID*)
   FBLKSIZE = 512;       (*STANDARD DISK BLOCK LENGTH*)
   DIRBLK = 2;           (*DISK ADDR OF DIRECTORY*)
   MAXSEG = 15;          (*MAX CODE SEGMENT NUMBER*)

TYPE

   DATEREC = PACKED RECORD
      MONTH: 0..12;               (*0 IMPLIES DATE NOT MEANINGFUL*)
      DAY: 0..31;                 (*DAY OF MONTH*)
      YEAR: 0..100                (*100 IS TEMP DISK FLAG*)
      END (*DATEREC*) ;

   UNITNUM = 0..MAXUNIT;
   VID = STRING[VIDLENG];

   DIRRANGE = 0..MAXDIR;
   TID = STRING[TIDLENG];

   FILEKIND = (UNTYPEDFILE,XDSKFILE,CODEFILE,TEXTFILE,
      INFOFILE,DATAFILE,GRAFFILE,FOTOFILE,SECUREDIR);

   DIRENTRY = RECORD
      DFIRSTBLK: INTEGER;                (*FIRST PHYSICAL DISK ADDR*)
      DLASTBLK: INTEGER;                 (*POINTS AT BLOCK FOLLOWING*)
      CASE DFKIND: FILEKIND OF
         SECUREDIR,
         UNTYPEDFILE:                    (*ONLY IN DIR[0]...VOLUME INFO*)
           (DVID: VID;                   (*NAME OF DISK VOLUME*)
            DEOVBLK: INTEGER;            (*LASTBLK OF VOLUME*)
            DNUMFILES: DIRRANGE;         (*NUM FILES IN DIR*)
            DLOADTIME: INTEGER;          (*TIME OF LAST ACCESS*)
            DLASTBOOT: DATEREC);         (*MOST RECENT DATE SETTING*)
         XDSKFILE,CODEFILE,TEXTFILE,INFOFILE,
         DATAFILE,GRAFFILE,FOTOFILE:
           (DTID: TID;                   (*TITLE OF FILE*)
            DLASTBYTE: 1..FBLKSIZE;      (*NUM BYTES IN LAST BLOCK*)
            DACCESS: DATEREC)            (*LAST MODIFICATION DATE*)
      END (*DIRENTRY*) ;

   SEGRANGE = 0..MAXSEG;
   SEGDESC = RECORD
      DISKADDR: INTEGER;      (* REL BLK IN CODE...ABS IN SYSCOM*)
      CODELENG: INTEGER       (*# BYTES TO READ IN*)
      END; (*SEGDESC*)

   vsegrange = 0..15; ( virtual segment number )

   ( segment table in syscomrec and in this program )
   segtabtype = array [segrange] of record
      codeunit : unitnum;
      codedesc : segdesc;
      end;

var
   alias : record case boolean of ( allow manual setup of ↑syscomrec )
      true : (i : integer);
      false: (p : ↑segtabtype);
      end;
   disk : file; ( file to read directory and seg tables from )
   ( global in which to store processed seg table for later use )
   vsegs : segtabtype;

segment procedure virtual;
   begin
   writeln('I am the REAL segment procedure 10.');
   end; ( virtual -- dummy )

procedure dovirtual(vsegnum : vsegrange);
   begin
   if (vsegs[vsegnum].codeunit = 0)  then begin
      writeln('Attempt to execute unlinked virtual seg number ',vsegnum,'.');
      exit(dovirtual); ( not linked )
      end;
   ( load segment register 10 with segment data from i'th segment procedure )
   alias.p↑[10] := vsegs[vsegnum];
```

```
   ( read in segment table )
   firstblk := directory[i].dfirstblk;
   unitread(unum,lsegtable,sizeof(lsegtable),firstblk);
   ( enter data from segment 10 in found file into v seg table )
   vsegs[vsegnum].codeunit := unum;
   vsegs[vsegnum].codedesc.codelens := lsegtable[10].codelens;
   vsegs[vsegnum].codedesc.diskaddr := lsegtable[10].diskaddr + firstblk;
   writeln('Finished with association of file ',fname,'.');
   end; ( virtlink )

begin
( initialize vseg table so that all units are zero (vseg undefined) )
for i := 0 to 15 do vsegs[i].codeunit := 0;
( set up pointer to syscomrec's seg table located by 'find' program )
alias.i := 718 + 96; ( 96 bytes in syscom rec before seg tables )
( read directory into memory )
unitread(unum,directory,sizeof(directory),dirblk,0);
numfiles := directory[0].dnumfiles;
writeln('Directory of unit ',unum,' read in.');
( link file names with v seg numbers )
virtlink('FAKE0.CODE',0);
virtlink('FAKE1.CODE',1);
virtlink('FAKE2.CODE',2);
virtlink('FAKE3.CODE',3);
virtlink('FAKE4.CODE',4);
virtlink('FAKE5.CODE',5);
virtlink('FAKE6.CODE',6);
virtlink('FAKE7.CODE',7);
end; ( virtinit )

begin ( main )
virtinit;
dovirtual(0);
dovirtual(1);
dovirtual(2);
dovirtual(3);
dovirtual(4);
dovirtual(5);
dovirtual(6);
dovirtual(7);
end.


program fake0;

   segment procedure fake0also;
      begin
      writeln('I think that I am virtual segment procedure 0.');
      end;

   begin ( main )
   end.
   virtual;
   end; ( dovirtual )

procedure virtinit;
( initialize tables for calling v segment procedures )
   const
      unum = 5; ( unit number where v segment procs are found )
   var
      i : integer; ( temp var )
      directory : array [dirrange] of direntry; ( holds disk directory )
      numfiles : dirrange; ( number of files in disk dir )

   procedure virtlink(fname : string; vsegnum : vsegrange);
      var
         i : integer;
         done : boolean;
         firstblk : integer; ( blk number of first blk of fake seg code file )
         lsegtable : array [segrange] of segdesc; ( holds seg tbl from disk )
      begin
      ( find file in directory )
      i := 0; done := false;
      while (i <= numfiles) and not done do begin
         i := i + 1;
         if (directory[i].dtid = fname) then done := true;
         end;
      if not done then begin
         writeln('unable to find file ',fname,'.');
         exit(virtlink);
         end;
      writeln('File ',fname,' found.');
```

Little-Ada* (Part II)

by Ralph E. Kenyon, Jr.

Little-Ada

Little-Ada version L/1 consists of a language design, a target machine design for intermediate code and a compiler to compile the language to the target machine code. The target machine is called the L-Machine.

Adding a system-dependent interpreter produces the specific L/1 implementation (in this case: Little-Ada version PolyMorphic L/1).

The Little-Ada design goal is to implement a subset of Ada, consistent with full Ada, which can support the basic Ada language structure and syntax. The purpose of the design is to produce a transportable compiler which can be implemented on various computers from the smallest to the largest, with an eye to future expansion to include more and more features of full Ada. The intended use of the L/1 version of Little-Ada was to support education in compiler principles and implementation methods in a classroom environment.

Little-Ada implements a minimal subset of Ada that preserves the block structuring and syntax forms in full Ada. Ada has been described as, in part, a block-structured language with strong typing that is Pascal-based. Little-Ada implements one aggregate (ARRAY), two objects (variable and constant), three control structures (subprogram, if-then-else, and loop-exit), relational and arithmetic operators and an assignment statement. Arithmetic operators include *, /, MOD, +, and -. All control structures in Ada and in Little-Ada conform to a general arrangement called a "comb." Examples include:

Program declaration	LOOP - EXIT	IF - THEN - ELSE
```/-- PROCEDURE .. IS```	```/-- LOOP```	```/-- IF .. THEN```
```¦ ..```	```¦ ..```	```¦ ..```
```¦-- BEGIN```	```¦-- EXIT```	```¦-- ELSIF .. THEN```
```¦ ..```	```¦ ..```	```¦ ..```
```\-- END;```	```\-- END LOOP;```	```¦-- ELSE```
		```¦ ..```
		```\-- ENDIF;```

Ralph E. Kenyon, Jr., 1686 West Main Rd., Portsmouth, RI 02871.
*Ada is a trademark of the U.S. Department of Defense (Ada Joint Program Office).

Reserved Words

Little-Ada has eighteen reserved words:

ARRAY	ELSE	EXIT	LOOP	OF	THEN
BEGIN	ELSIF	INTEGER	MOD	PRAGMA	TYPE
CONSTANT	END	IS	NULL	PROCEDURE	WHEN

There are eighteen special symbols in Little-Ada.

```
+  -    > =   :=   *   /
<  >    < =   ..   =   :
(  )    /=    --   ;   _
```

a. Six condition test symbols:

```
=  /=  <   < =  >   > =
```

b. Four arithmetic operators:

```
+  -  *  /   MOD
```

(MOD is included here for completeness.)

c. Assignment operator (becomes):

```
:=      (as in NEW := OLD + 1;)
```

d. Object type indicator:

```
:   (as in SIZE : INTEGER;
```

e. Range indicator (ellipsis read "to"):

```
..   (as in ARRAY (1..10) OF)
```

f. Comment (Terminated by end of line):

```
--
```

g. Statement terminator:

```
;
```

h. Parentheses:

```
( )
```

i. Underscore:

```
_ (ignored in numbers)
```

Declarations

In Little-Ada, all program units must be declared. Program units and sub programs are declared in exactly the same way. The following is an example.

```
PROCEDURE NAME IS
  -- declare part
BEGIN
  -- sequence of statements
END;
```

Little Ada, continued...

The declare part must declare all objects not already declared that are subsequently referenced in the procedure. Little-Ada differs from full Ada in that Little-Ada (L/1) does not implement parameters, and locally declared variables hide global variables of the same name. Parameter passing can be achieved in Little-Ada in a manner that is simpler to implement, but does not provide for safer programming later. In Little-Ada, procedures are allowed access to all variables in the same scope (not hidden).

An ARRAY type may be declared with the following syntax:

```
TYPE NEW←TYPE IS ARRAY (LOW..HI) OF BASE←TYPE;
```

Examples:

```
TYPE VECTOR IS ARRAY (1..10) OF INTEGER;
TYPE MATRIX IS ARRAY (1..10) OF VECTOR;
```

Variables and constants are declared as follows:

```
VARIABLE←NAME : TYPE←NAME;
CONSTANT←NAME : CONSTANT INTEGER := VALUE;
```

Examples:

```
SIZE : INTEGER;
HIGHT : INTEGER := 10;
LINES←PER←PAGE : CONSTANT INTEGER := 66;
```

In Ada and Little-Ada, identifiers and numbers may contain single embedded underscores for readability. Examples:

```
LOW←IN←CH;  22←350
```

The underscore has no effect on numbers, but is part of identifiers. An identifier starts with a letter, but may contain digits and letters with single embedded underscores. Little-Ada only distinguishes the first ten characters in an identifier and limits numbers up to 32←767. Also, in Little-Ada only upper case may be used in identifiers and keywords. Comments may use both upper and lower case. Here Little-Ada differs from full Ada in that full Ada accepts lower case, but treats lower case as upper case in identifiers and reserved words.

Statements

Little-Ada implements only one of the Ada LOOP structures:

```
LOOP
--
EXIT;
END LOOP;
```

EXIT statements must be used to exit from a Little-Ada LOOP statement. The syntax is as follows:

```
EXIT;
-- or
EXIT WHEN condition;
```

Examples:

```
EXIT WHEN LOW←CHAR = 13; -- end of line
IF LOW←CHAR = 13 THEN EXIT; END IF;
-- (same effect)
```

Comparing loop structures in Basic, Fortran, Little-Ada and Pascal:

Little-Ada		Basic
LOOP		10 REM
		20 REM statements
-- statements		30 IF condition GOTO 60
EXIT WHEN CONDITION;		40 REM more statements
-- statements		50 GOTO 10
END LOOP;		60 REM

```
           Fortran          :        Pascal
------------------------------+--------------------------
                              :     LABEL 10,20;
                              :
 10        CONTINUE           :    10 NIL;
 c         statements         :          { statements; }
           IF condition GOTO 20 :    IF condition THEN
 c         more statements    :             GOTO DONE;
           GOTO 10            :          { more statements; }
 20        CONTINUE           :             GOTO 10;
                              :    20 NIL;
```

```
   Other Pascal loops          : Little-Ada
---------------------------------+------------------------
        -          : REPEAT      : LOOP
        -          :   BEGIN     :   --
        -          :   { statements; } :   -- statements
        -          :   END;      :   --
WHILE NOT condition : UNTIL condition; : EXIT WHEN cond;
  BEGIN            :      -       :   --
    { statements; } :      -      :   -- statements
  END;             :      -       :   --
END;               :      -       : END LOOP;
```

Full Ada has a FOR K IN range LOOP .. END LOOP; (where "range" is from LOWER←BOUND to UPPER←BOUND), which can be accomplished in Little-Ada as follows:

```
I : INTEGER;

I := LOWER←BOUND;
LOOP
EXIT WHEN I>UPPER←BOUND;
   -- statements
   I := I + 1;
END LOOP;
```

Corresponding structures are:

```
   Basic         :    Fortran      :   Pascal
------------------+-----------------+----------------
10 FOR I=1 TO N   :    DO 20 I=1,N  : FOR I := 1 TO N DO
20 REM statements :  C  statements  :   BEGIN
30 NEXT I         :  20 CONTINUE    :     { statements; }
                  :                 :   END;
                  :                 : END;
```

Conditional control structure in Little-Ada and in full Ada uses the IF statement. The syntax is:

```
IF condition THEN statements;
ELSIF condition THEN statements;
ELSE statements;
END IF;
```

ELSEIF and ELSE are optional. Examples:

```
IF SCORE > 95              IF DOOR = OPEN
   THEN MARK := A;            THEN CLOSE;
ELSIF SCORE > 85           END IF;
   THEN MARK := B;
ELSIF SCORE > 75
   THEN MARK := C;         IF SEX = MALE
ELSIF SCORE > 65           THEN
   THEN MARK := D;            COLOR := BLUE;
ELSE                       ELSE
   MARK := F;                 COLOR := PINK;
END IF;                    END IF;
```

The primary statement is the assignment statement.

```
RESULT := EXPRESSION;
```

Numerical computations include +, -, *,/ and MOD, and are all 16-bit signed integers. Multiplicative operators *, /, and MOD take presidence over additive operators + and -. Multiplicative operations should be parenthetically grouped to prevent ambiguity. A * B / D is ambiguous and should be written (A * B) D, or A * (B / D).

L/1 Machine

The L/1 version of the L - Machine is a 16-bit stack-oriented processor with a single 8-bit I/O port and a simple instruction set consisting of six primary and eighteen secondary instructions.

L/1 Compiler

The compiler implements only one data type—16-bit integers. INTEGER is a built-in type. One built-in variable, LOW←CHAR, is used to load the I/O port. Two build-in procedures, LOW←IN←CHAR and LOW←OUT←CH input and output from LOW←CHAR through the I/O port. While Ada allows use of upper and lower case, Little-Ada only recognizes upper case in program identifiers, while comments may include lower case.

Additional limitations of the compiler include a limit of 2000 bytes of code and fifty symbols. Because of the recursive descent technique used by the compiler, the fifty symbols mean a maximum of fifty symbols at a given moment during compilation. Space for local names within a procedure is reclaimed when that procedure has completed compiling. This allows more than fifty total symbols, provided the current stage is never more than fifty. Procedures may be nested to a maximum of fifteen levels (LOW←IN←CHAR and LOW←OUT←CH are one level below the main program). Input lines may be a maximum of 81 characters (including the CRLF). Each object declaration is limited to a maximum of seventeen identifiers per statement.

There are some more subtle limitations to the compiler. It does not make any run-time checks of array limits. The compiler will also not detect some kinds of syntax errors. Initializing a CONSTANT or a variable will not detect an error if an identifier is used instead of a number. Assignments of an array will not assign the entire array; only the first element is assigned.

Sample Programs

The following example is an illustration of a simple program to output a carriage return.

```
PROCEDURE CR IS
  --
BEGIN -- CR
  LOW<-CHAR := 13;
  LOW<-OUT<-CH; -- output LOW<-CHAR
END; -- CR
```

The following example illustrates the use of this program as a subprogram in another program called HI. Subprograms are invoked by mentioning their names.

```
PROCEDURE HI IS
  PROCEDURE CR IS
  BEGIN -- CR
    LOW<-CHAR := 13;
    LOW<-OUT<-CH; -- output LOW<-CHAR
  END; -- CR
BEGIN -- HI
  LOW<-CHAR := 72; -- H
  LOW<-OUT<-CH; -- output "H"
  LOW<-CHAR := 105; -- i
  LOW<-OUT<-CH; -- output "i"
    CR; -- output a CR too
  END; -- HI
```

The following is a more significant program. This program illustrates pseudo parameter passing as well as programming style. It prints out a list in columns of numbers from zero to fourteen which includes the number, the powers of two, the squares of the number, the sum of the numbers, etc.

```
PROCEDURE EXAMPLE IS
  N : INTEGER; -- N is the global IN parameter
  END<-OF<-LINE : CONSTANT INTEGER := 13;
  W,X,Y,Z,J : INTEGER;
  --
  PROCEDURE PRINT<-N IS
    TYPE NUMBER IS ARRAY (1..6) OF INTEGER;
    -- ASCII constants
    ZERO : CONSTANT INTEGER := 48;
    BLANK : CONSTANT INTEGER := 32;
    MINUS : CONSTANT INTEGER := 45;
    -- variables
    DIGIT : NUMBER;
    X,INDEX : INTEGER; -- X is work copy
  BEGIN -- PRINT<-N
    -- Copy global parameter into local IN parameter X
    X := N;
    -- first, check sign
    IF X < 0
    THEN -- sign goes in the '6th' digit position
      DIGIT (6) := MINUS - ZERO; -- we add ZERO on output
      X := -1 * X; -- MOD only works on positive numbes
    ELSE
      DIGIT (6) := BLANK - ZERO;
    END IF;
    --
    INDEX := 0;
    LOOP -- now we put each digit in its array slot
      INDEX := INDEX + 1;
      DIGIT (INDEX) := X MOD 10;
      X := X / 10;
    EXIT WHEN INDEX = 5;
    END LOOP;
    --
    LOOP -- now we blank leading zeros and move sign
    EXIT WHEN INDEX = 1;
      IF DIGIT (INDEX) = 0
      THEN -- move sign
        DIGIT (INDEX) := DIGIT (INDEX + 1);
        -- and blank previous position
        DIGIT (INDEX + 1) := BLANK - ZERO;
      ELSE
      EXIT; -- if the digit wasn't zero we want to quit
      END IF;
      INDEX := INDEX - 1;
    END LOOP;
    --
    LOW<-CHAR := BLANK;
    LOW<-OUT<-CH; -- Two leading blanks
    LOW<-OUT<-CH; -- on each number
    INDEX := 6;
    LOOP -- now we print out all six digits
    EXIT WHEN INDEX = 0;
      LOW<-CHAR := DIGIT (INDEX) + ZERO;
      LOW<-OUT<-CH;
      INDEX := INDEX - 1;
    END LOOP;
  END; -- PRINT<-N

  --
  BEGIN -- EXAMPLE
    W := 0;
    J := 0;
    Y := 0;
    Z := 1;
    LOOP
      N := J;            -- INDEX
      PRINT<-N;
      N := Z;            -- 2↑INDEX
      PRINT<-N;
      N := J * J;        -- INDEX↑2
      PRINT<-N;
      N := Y;            -- SUM J=0 TO N
      PRINT<-N;
      N := W;            -- SUM J↑2 TO N
      PRINT<-N;
      N := Z - W;        -- 2↑N - SUM N↑2
      PRINT<-N;
      LOW<-CHAR := END<-OF<-LINE;
      LOW<-OUT<-CH;
    EXIT WHEN J = 14;
      Z := Z * 2;
      J := J + 1;
      W := W + J * J;
      Y := Y + J;
    END LOOP;
  END; -- EXAMPLE
```

Program EXAMPLE Output

0	1	0	0	0	1
1	2	1	1	1	1
2	4	4	3	5	-1
3	8	9	6	14	-6
4	16	16	10	30	-14
5	32	25	15	55	-23
6	64	36	21	91	-27
7	128	49	28	140	-12
8	256	64	36	204	52
9	512	81	45	285	227
10	1024	100	55	385	639
11	2048	121	66	506	1542
12	4096	144	78	650	3446
13	8192	169	91	819	7373
14	16384	196	105	1015	15369

```
                    BLOCK                                    IDENTIFIER

    :...........................<;>...[type]..........            ...........[letter]...
    :/                                              \            :        ...[digit].../   \...........
    :\....<TYPE>..[identifier]...<IS>..........       \           :/                      \
    :                                        \        :           :\....<underscore>.../:
    :           .........<ARRAY>.............../                   \................../
    :          /
    :   \..<()....[number]...<,>...[number]...<)>...<OF>./

    :.............................................<;>.....
    :/                                                   \
    :\...[identifier]..<:>....................[type]...../
    :   /              \      \                 \        :
    :   \......<,>...../       \..<CONSTANT>../           :            CONDITION
    :          .....................................
    :          \...<:=>.......[number]............/         ....[expression].............<=>.........../...[expression]....
    :/                                             \                               \        /
    :\.....<PROCEDURE>..[identifier]...<IS>...[block]....../                         :\..</=>.../:
    :                                                                               :\...<<>.../:
    \...<BEGIN>...[sequence of statements]....<END>...<;>...                        :\..<<=>.../:
                                                                                    :\...<>>.../:
                                                                                     \...<>=>../


                    NAME                                                        EXPRESSION

    .................................[identifier]..........................      ...........................[term]......
    \                                                      /                     \                        /         \
     \...[name]........<()>.....[expression]....<)>...../                         :    ....<+>..../        :
                                                                                  :    \.....<->..../     :
                                                                                   \.....................


                    NUMBER                                                         TERM

    .........[digit].............                           ................../[factor].................
    /                    \                                                   :
    :\...<underscore>.../:                                                   :\.....<*>...../:
    :                    :                                                   :\.....</>...../:
    \.................../                                                     \.....<MOD>..../


                    PROGRAM
    ......<PROCEDURE>....[identifier]....<IS>....[block].....                       FACTOR

            SEQUENCE OF STATEMENTS                            ......................[number].....................
                                                             \                                        /
            \...........                 ............/        :\................[name].............../:
                        \             /                       :\....<()....[expression]....<)>....../
    ....................../...........\...<;>.................
    /
    :\......[variable name].....<:=>.....[expression]........./
    :\.........[subprocedure identifier]..................../
    :\....<LOOP>...[sequence of statements]......<END>..<LOOP>.../:
    :\...<EXIT>........<WHEN>....[condition].................
    :          \
    :           \...................................
    :\.<IF>....[condition]..<THEN>..[sequence of statements]..
    :  /                                              \
    :  \....<ELSIF>....................................../:
    :          /
    :           \..<ELSE>..[sequence of statements]...
    :  ...................
    : :
    :  \....<END>..<IF>...................................
    \.................<NULL>............................./
```

                                                    Little-Ada (L/1) Syntax Diagrams

Block	Identifier
Name	Condition
Number	Expression
Program	Term
Sequence of Statements	Factor

Table V: *Little-Ada. L/1 Compiler Error Listing.*

Error	Interpretation		Error	Interpretation
001	Program too big! Code overflow!		031	Missing "LOOP" on END statement!
002	Compiler error! Patch not a branch!		032	No EXIT from LOOP!
003	Undefined base type!		033	EXIT statement begin error!
004	Inappropriate base type!		035	Undefined identifier!
005	Too many names! Symbol table overflow!		036	Not variable nor subprogram!
006	Number too large!		037	EXIT not allowed here!
007	Illegal character > "Z"!		038	Missing semicolon at statement end!
008	Illegal character > "9" and < "A"!		039	Compiler error! TYPE token lost!
009	Illegal character < "0"!		040	Missing type name!
010	Number too large!		041	Missing IS!
011	Undefined identifier!		042	Missing "ARRAY" in type declaration!
012	Array of invalid type!		043	Missing left parenthesis!
013	Missing right parenthesis!		044	Missing lower bound!
014	Illegal use of identifier!		045	Missing ".."!
015	Missing right parenthesis!		046	Missing upper bound!
016	Illegal factor beginning!		047	Missing right parenthesis!
017	Illegal relational operator!		048	Missing OF!
018	Illegal subscripting!		049	Missing base type name!
019	Missing right parenthesis!		050	Too many identifiers in declaration!
020	Missing ":="!		051	Missing semicolon!
021	IF statement begin error!		052	Missing semicolon after unit body!
022	Compiler error! THEN token lost!		053	Level too deep!
023	Missing THEN!		054	Missing BEGIN on unprefixed block!
024	Missing END!		055	Missing END on unprefixed block!
025	Compiler error! IF token lost!		065	Subprogram begin error!
026	LOOP statement begin error!		066	Missing subprogram name!
030	Missing "END" at end of LOOP statement!		067	Missing IS in procedure declaration!
			068	Semicolon expected after last "END"!

Little-Ada.L1 Test Program

This is an example of a Little-Ada.L1 test program. The output here is exactly as it appeared running the L/1 Interpreter on the block of code developed by Dr. Mathis. The block of code is the transportable compiler. In the listing, the code index is followed by a copy of the input line of text. Errors will show up in the listing if any exist (none shown here). See Table V for a complete list of errors generated by the Little-Ada compiler. The test program listing is followed by an ASCII block of the output code (which is the L/1 object code for this program).

```
CODE
INDEX--SOURCE CODE--
0006 PROCEDURE TEST IS
0006 TYPE NUMBER IS ARRAY(1..5) OF INTEGER;
0006 VAL:NUMBER; I,X,Y:INTEGER;
0014  -- A program to print the number 00256
0014 BEGIN
0016    X:=256;
0022    -- This routine puts digits into array value
0022    I:=1; LOOP EXIT WHEN I>5;
0038    Y:=X/10; VAL(I):=X-Y*10; X:=Y;
0078    I:=I+1; END LOOP;
0090    -- This routine prints out the digits top first
0090    I:=5; LOOP EXIT WHEN I=0;
0106    LOW←CHAR:=VAL(I)+48; LOW←OUT←CH;
0127    I:=I-1; END LOOP;
0139    --
0139 END;
**
0006 908C 918C A58F A18F A18F A18F 0010 C000 09B1 008E C000
0AA1 8EC0 000A 8DA5 8A87 4026 005A C000 08C0 0009 8DAA 848E
C000 03C0 000A 8DA1 8281 C000 098D C000 088D AA83 828E C000
09C0 0008 8D8E C000 0AC0 000A 8DA1 818E 001B C000 0AA5 8EC0
000A 8DA0 8887 406A 008B C100 00C0 0003 C000 0A8D A182 818D
B030 818E D100 04C0 000A C000 0A8D A182 8E00 5F8C 0000 0000
```

The above block of code was "hand disassembled" to produce the L/1 assembly source listing shown below. The comments show the relationship of the Little-Ada.L1 test program to the generated code.

ADDR CODE	LABEL	OPCODE	COMMENT
		ORG 0	
0000 0006		br Main	

;This is a subroutine for the built
;in routine to input a character

ADDR CODE	LABEL	OPCODE	COMMENT
0002 90	IN←CH	inb	;LOW←IN←CH
0003 8C		ret	

;This is a subroutine for the built
;in routine to output a character

ADDR CODE	LABEL	OPCODE	COMMENT
0004 91	OUT←CH	outb	;LOW←OUT←CH
0005 8C		ret	

;Code generated for TEST.AD starts
;here. Initially, space is reserved
;on the stack for the declared
;variables and constants. Each is
;done separately, as it is
;processed by the parser.

ADDR CODE	LABEL	OPCODE		COMMENT
0006 A5	Main	lic 5		;VAL
	VAL	EQU	3	
0007 8F		inc		
0008 A1		lic 1		;Y
	Y	EQU	8	
0009 8F		inc		
000A A1		lic 1		;X
	X	EQU	9	
000B 8F		inc		
000C A1		lic 1		;I
	I	EQU	0AH	
000D 8F		inc		

;After all declarations are processed,
;the program branches past any
;subprocedures to the beginning of
;the main part. Here, there are no
;subprocedures so the branch is
;to the next address.

ADDR CODE	LABEL	OPCODE	COMMENT
000E 0010		br Start	

;The assignment statement X := 256;

ADDR CODE	LABEL	OPCODE	COMMENT
0010 C0 0009	Start	lad 0,X	
0013 B1 00		lic 256	
0015 8E		sto	;X:=256

;The assignment statement I := 1;

ADDR CODE	LABEL	OPCODE	COMMENT
0016 C0 000A		lad 0,I	
0019 A1		lic 1	
001A 8E		sto	;I:=1

;The LOOP statement does nothing
;except establish the label LOOP1
;for the backward branch at the
;corresponding END LOOP;

;The EXIT WHEN condtion statement
;caused the condition code
;generation first and the
;conditonal branch for exit.

```
ADDR CODE      LABEL    OPCODE        COMMENT

001B C0 000A   LOOP1    lad  0,I
001E 8D                 rav
001F A5                 lic  5
0020 8A                 setst         ;WHEN I>5
0021 87                 not
0022 4026               bnz  BLOCK1
0024 005A               br   BLOCK2   ;EXIT

0026 C0 0008   BLOCK1   lad  0,Y
0029 C0 0009            lad  0,X
002C 8D                 rav
002D AA                 lic  10
002E 84                 div
002F 8E                 sto           ;Y:=X/10
0030 C0 0003            lad  0,VAL
0033 C0 000A            lad  0,I
0036 8D                 rav
0037 A1                 lic  1
0038 82                 sub
0039 81                 add           ;VAL(I)=VAL+(I-1)
003A C0 0009            lad  0,X
003D 8D                 rav
003E C0 0008            lad  0,Y
0041 8D                 rav
0042 AA                 lic  10
0043 83                 mul           ;10*Y
0044 82                 sub           ;-X
0045 8E                 sto           ;VAL(I)
0046 C0 0009            lad  0,X
0049 C0 0008            lad  0,Y
004C 8D                 rav
004D 8E                 sto           ;X:=Y
004E C0 000A            lad  0,I
0051 C0 000A            lad  0,I
0054 8D                 rav
0055 A1                 lic  1
0056 81                 add
0057 8E                 sto           ;I:=I+1
0058 001B               br   LOOP1
005A C0 000A   BLOCK2   lad  0,I
005D A5                 lic  5
005E 8E                 sto           ;I:=5
005F C0 000A   LOOP2    lad  0,I
0062 8D                 rav
0063 A0                 lic  0
0064 88                 sete          ;WHEN I=0
0065 87                 not
0066 406A               bnz  BLOCK3
0068 008B               br   End      ;EXIT

006A C1 0000   BLOCK3   lad  1,LOW+CHAR
006D C0 0003            lad  0,VAL
0070 C0 000A            lad  0,I
0073 8D                 rav
0074 A1                 lic  1
0075 82                 sub
0076 81                 add
0077 8D                 rav           ;VAL(I)=VAL+(I-1)
0078 B030               lic  '0'
007A 81                 add
007B 8E                 sto           ;LOW+CHAR:=VAL(I)+'0'
007C D1 0004            call 1,OUT+CH
007F C0 000A            lad  0,I
0082 C0 000A            lad  0,I
0085 8D                 rav
0086 A1                 lic  1
0087 82                 sub
0088 8E                 sto           ;I:=I-1
0089 005F               br   LOOP2

008B 8C        End      ret
                        END
```

Summary

Emulator programs can be written for machines that do not exist. These non-exsistant machines are useful in the design of hardware machines, as well as in the design of high level languages. One such emulator on the 8080A in the PolyMorphic 8813 is being used in the development of a compiler for a subset of the new DOD language Ada called "Little-Ada." Ada is a strongly typed block structured language being designed for use in DOD embedded computer systems. The emulated machine is called the L - Machine and has been described in detail. This emulation also shows how a stack-oriented machine of radically different characteristics than the 8080A can be implemented on an 8080A. ∎

The Tarbell Double Density Disk Interface

by Fred Greeb

You want to use a DMA controller with dynamic memory? Good luck! That was the typical reaction when I said that I was considering ordering the Tarbell Double Density disk controller board. For those of you who don't know, a DMA controller performs data transfer directly to and from the computer memory, without any processor intervention required. It essentially turns off the main system processor, and the processor timing, and generates its own timing signals for data transfer. Dynamic memory boards, such as the SD Sales Expandoram which I have in my system, tend to be very fussy and require that the timing signals be just right for proper operation. That's why I got the reaction that I did. I didn't pay any attention to these comments, and went right ahead and ordered the controller. This article describes my experience in getting the board running in my system.

Selection of the Controller

There were several reasons why I selected the Tarbell controller. I plan to expand my computer in the future to include multi-user software, and wanted a disk controller which would not tie up the system during disk data transfer. If the controller requires the processor to handle the data, normal terminal input cannot be accomplished during disk accesses. This can be very annoying in a multi-user environment.

To meet this requirement, I needed either a DMA controller, or a controller with an on-board buffer for temporary storage of the disk data. The controllers with on-board memory buffers that I looked at all required that the controller memory occupy part of the main memory space, usually somewhere in the top 16K of the normal 64K addressing range of an 8-bit microprocessor. Since I wanted to retain the capability to expand the memory beyond 64K, with extended addressing rather than bank select, this was not an acceptable option.

The Tarbell controller implements full 24-bit extended addressing during DMA data transfers, utilizing the additional address lines defined in the IEEE-696/S-100 bus

specification. This feature would be very useful if and when I expand the system to either multi-user software or 16-bit processor.

Finally, I have been running the Tarbell single density disk controller for over three years, and have not had any problems with it. The single density controller has proven to be a very reliable unit. This fact gave me confidence in Tarbell's capability to design and manufacture a double density interface. Since I already had two 8" disk drives, the fact that the controller does not contain the option to use 5 1/4" mini drives was of no concern.

The Controller Arrives

I ordered the controller board from one of the mail order computer suppliers which regularly advertises discount prices in most of the computer magazines. I ordered an assembled and tested unit, since I was told that Tarbell no longer sells the controller in kit form. The unit I received was marked Revision G. Hopefully, I thought, with that many revisions, all of the bugs will have been removed from the design.

The board is rather densely packed with about fifty integrated circuits, all socketed, and numerous discrete components to handle the disk to computer interface. Most of the jumper options that were available on the Tarbell single density controller have been removed. The double density unit is defined as Shugart-compatible only, and does not contain jumper options to configure it for other drives. Shugart-compatible does not mean that you have to use Shugart drives. It only requires that your drives utilize the same interface signals on the same connector pins as the Shugart drives. The majority of the drives currently available meet this requirement.

The documentation which arrived with the controller consisted of a manual and a disk. The manual is adequate, but not overly detailed in the technical description section. There is a very good section on getting the controller running, especially if you have been running the single density controller as I had been. The manual contains over two pages covering "what to do if it doesn't work," and listing symptoms and possible cures. It also contains a list of compatible processor boards, memory boards,

Fred Greeb, 1915 S. Cape Way, Denver, CO 80227.

and I/O boards which have been tested with the controller. The last half of the manual contains specification sheets for all of the integrated circuits used in the controller, including the 8257 DMA controller chip, and the 1793 floppy disk controller chip. A complete schematic of the controller takes up the last two pages of the manual.

The disk included with the interface is labelled "Public Domain Disk #2," with a copyright notice which seems to conflict with the public domain statement. The content of the disk is extremely helpful. Included on the disk are all the I/O drivers necessary to make the double density controller operate with CP/M, and several disk formatting and testing programs. If all manufacturers would supply programs required to make their hardware operate, in machine readable form, life would be easier for the microcomputer enthusiast.

Bringing Up the Controller

Bringing up the controller was very easy, since I already had a working disk system. I merely copied the BIOS file from the Public Domain disk, edited it to match the requirements of my terminal, and assembled the edited file. After merging the assembled file with my CP/M operating system, I had a disk ready to run with the new controller.

I installed the new controller board, powered up the computer, and tried to boot the newly created disk. It didn't work! Not to be discouraged, I tried booting the bootstrap ROM contained on the controller board. This time everything worked perfectly. The disk loaded, and everything appeared to be operating normally. The only problem was with the bootstrap ROM. This is apparently a common problem, and is dependent on the processor board (and motherboard) which you are using. The manual discusses this problem, and gives the probable cures. In my case, changing a capacitor from 100 pf to 370 pf solved the problem. The controller was working fine, including the bootstrap operation.

DMA Operation

Up to this point, I had the controller working in programmed I/O mode, which requires active processor participation. In this mode, the operation of the double density interface is very similar to the older single density interface. The next step was to implement the DMA mode, which would allow me to run double density. Programmed I/O control requires a 4MHz system for double density, and my old Altair computer just won't run that fast.

The only change necessary to try the DMA mode was to change one equate statement in the BIOS file contained on the public domain disk. I edited the file, reassembled it, and placed a DMA version of the operating system on my disk. I hit the reset button to boot up the DMA mode operating system, hoping that the controller and my dynamic memory would co-exist peacefully. It was with much relief that I read the sign-on message from the operating system, indicating that everything had worked properly. I was running DMA mode, single density, and my memory board appeared perfectly happy.

Then I tried copying a file from one disk to another. The copy contained garbage! This was the first operation I had performed which required a DMA read operation

Response from Tarbell Electronics:

We have never claimed to be completely compatible with the proposed IEEE-696 standard, because we found it difficult to determine how close we were to the standard. We have, however, made it our policy to design our boards as close to the standard as possible without compromising features.

Revision G of our board was designed specifically to work with Ithaca InterSystems 64K dynamic memory board, and will still work with the Measurement System 64K dynamic memory board. We have too little experience with other dynamic boards to say with which it will or will not work. We are glad to hear, of course, that with a small modification, it worked with the SD Sales Expand-oram.

DMA bus arbitration was not included for three reasons:

1) We would have had to take out one of the other features, for example the extended address-bus feature, with the parts available at the time. There are a lot of parts on the board.

2) It wasn't clear to us from the standard whether the arbitration was a required feature to meet the standard.

3) We felt that most installations wouldn't require that more than one DMA device transfer data at a time. Note that it is possible to have our interface on the bus with another DMA interface, as long as both aren't trying to transfer data at the same time. Since the most likely candidate for the other device would be a hard disk interface, it would most likely require data transfers at too high a rate to allow our interface to transfer concurrently anyway.

We welcome any other ideas from your readers about how any of our products can be improved.

Don Tarbell
owner

from memory. Apparently I could do a DMA write to memory, to load the operating system from disk to memory, but could not do a DMA read to go from memory to disk.

This problem was not mentioned in the manual, so I had no clues as to what was happening. I studied the memory timing specifications to determine what was required by the memory board. Then I determined the timing signals generated by the DMA controller circuits. It would have been most helpful at this point if timing diagrams had been included with the controller documentation.

Anyway, the two signals which were causing the problem were the pSYNC signal, which defines the start of a new bus cycle, and the sMEMR signal, which identifies a data transfer from memory. The memory board requires that both the pSYNC and sMEMR signals be active to initiate a memory read cycle. The DMA controller sets pSYNC active at the start of a DMA cycle, and then turns it off as soon as the sMEMR signal goes active. My memory board would never see both of these signals active at the same time, and therefore would never initiate a read cycle. Either the memory board or the disk controller would have to modified.

64   MICROSYSTEMS

I redesigned (on paper) the DMA control circuit and the memory read circuit, trying to correct this problem, but could not come up with an easy solution. Finally, I just lifted the pin which controlled pSYNC on the controller board, making the pSYNC signal true through the entire DMA cycle, and tried the copy operation again. This time everything worked fine. I could read and write disk files, and my memory board responded correctly. I am still running in this configuration. Not exactly IEEE spec, but it works.

Double Density Operation

With the DMA problem solved, I was ready to try double density operation. I formatted a disk using the disk format program supplied with the controller, and ran the disk test program. Errors and more errors! This problem was well documented and relatively easy to cure. As described in the manual, different disk drives require different pre-compensation, but I could follow the instructions provided in the manual to correct the problem. Just try different settings on the nine switches on the controller board until the errors go away. Various combinations of switch settings are listed in the manual and it is a straightforward, if somewhat time-consuming, process to configure the interface for the drives you are using. You change the switch settings, format and test a disk, and repeat the operation until the best combination is found. The switches are located at the top of the board and are easy to change.

After I had gone through this process I could operate in the double density mode without problems. The controller does seem to be a bit fussy as to the quality of the disk that I use. Some of my old single density disks will not format properly in the double density mode, but I have had no problems with certified double density disks.

Conclusion

At this point I have no major complaints with the double density interface. It is performing well, and I have over twice the disk space available compared to the older single density interface that I was using. I completely rewrote the BIOS file supplied with the controller, and achieved 596K bytes per disk, using 512 byte sectors.

My only minor complaint is in the area of IEEE compatibility. The controller is claimed to be compatible with the new IEEE specification, but only partially meets the requirements. In particular, the board operates in the DMA mode, but includes no DMA arbitration circuitry to handle multiple DMA devices on the bus. If you are going to install several DMA devices, this could lead to problems. I currently have no such plans, so I won't worry about it. Also, since no timing diagrams are included with the documentation, it is not easy to determine if the controller meets the IEEE timing specification.

The software supplied with the interface is one of the nicest features of the whole system. If I had to write all of that software from scratch, it would have slowed down the whole process considerably. Tarbell Electronics is to be congratulated for including the disk with the controller, even though I did not buy the CP/M operating system from them. ∎

Book Review

Osborne's CP/M User's Guide Reviewed

by Chris Terry

Osborne CP/M User's Guide, by Thom Hogan. Osborne/McGraw-Hill, Berkeley, CA 283 pp., $12.99. 1981.

This is a reference book for the CP/M operator, as distinct from the systems programmer. It is a thorough and very readable treatment of what the system does, how to use the built-in commands and utilities supplied with the package, and contains good information on high-level languages and application software to run under CP/M. Examples of usage are plentiful, clear, and well-commented. In these respects it is first-rate.

However, every silver lining has a cloud, as they say. Here the cloud is Mr. Hogan's restricted view of what a user is or wants. Surely, a CP/M user is anyone who installs CP/M in his system? Some users may not be technically oriented and quite willing to scream for an expert if the system so much as hiccups; these are well served by Mr. Hogan. But why does he choose to limit his teaching to this particular body of users? There are many others who are not experts but are anxious to learn something about the innards of the system they use, so that they can become, if not experts, at least competent to make simple modifications—much as a car owner will learn how to change a wheel in case he gets a flat, how to check oil, water, brake and transmission fluids, or how to look for and change a dying radiator hose. This kind of user will still find Mr. Hogan a good CP/M driving instructor, but must go to other sources (such as those in the bibliography) for instruction on preventive or corrective maintenance. I think this is a pity. If Mr. Hogan had handled the technical and interfacing aspects of CP/M with the same clarity and thoroughness that he gave to the operating aspects, we might have had the long-awaited definitive book on CP/M. It would have been very big, but also very understandable. Let us all hope that he will drop the other shoe—soon!

What The Book Contains

Chapter 1 is an introduction to operating systems in general and CP/M in particular. It gives the history of CP/M, speaks of the main versions of CP/M from Digital Research and how other vendors customize these, lists the manuals supplied with versions 1.4 and 2.x, and briefly discusses the function of CP/M in a microcomputer system. There is much useful information about the differences to be expected between versions sup-

plied by Digital Research and those supplied by other vendors (including "look-alikes" such as TP/M and CDOS). On the debit side, the section "Some Useful Terms" contained only "Bytes," "Tracks," and "Sectors." This is hardly adequate for what follows. There are twenty or so vital terms (such as "logged-in," "default," etc.) that should have been included even if it meant moving the glossary to an appendix.

Chapter 2 discusses the file-naming conventions and gives a detailed description of each of the built-in commands, including the control characters for editing the command line. On file naming, there are very clear explanations of how to use the "wild-card" characters "?" and "*", as well as explanations of why you can *not* use them in certain cases.

I particularly liked the remarks on the differences in how Versions 1.x and 2.x execute the commands, and the very clear definitions of the conventions used in this and other chapters to show precisely what the operator should type, and what the computer displays. Commands and computer responses are first shown single-spaced, as they would appear on the console, and then dissected with an explanatory comment on each part of each command or response. Additional paragraphs contain explanatory notes, or list the error messages that might appear and their probable causes. Excellent stuff. Chapters 2 through 4 and Chapter 6 all have this level of detail, and even a CP/M "expert" may learn a few things from these chapters.

Chapter 3, "CP/M Transient Commands," discusses what a program is, and shows that all commands other than the seven(six for Version 1.4)built-in commands are executed by programs resident on the disk and brought into memory by typing the program name as a command. The addition of parameters to commands is discussed in some detail. The rest of the chapter discusses the housekeeping utilities STAT, PIP, ED, DUMP, and the batch processing utilities SUBMIT and XSUB. It is not clear to me why DUMP is here—it would be more logical to put it in Chapter 4 (Assembly Language utilities), but that is a very minor quibble. Again, there are very detailed, accurate, and thorough discussions of the functions of these programs, how to invoke them, and what results to expect. The discussion of PIP, in particular, is the best I have seen and includes even exotica such as the "special" PIP devices NUL:, EOF:, INP:, OUT:, and PRN:.

Chapter 4, "Assembly Language Utilities," provides full and clear descriptions of ASM, DDT, and LOAD, and exactly how to use them. Once again, it is the explana-

Chris Terry, 324 E. 35th St., New York, NY 10016.

tory text and the examples that make this chapter valuable. Where the Digital Research manual curtly informs you that a Phase error results when "label does not have the same value on two subsequent passes through the program," this chapter notes: "The label changes value during assembly. If you must reassign a value, use the SET directive. Can also be caused by a duplicate label." A duplicate is indeed the most common cause. There are similar helpful comments on all of the ASM error messages. In this section on LOAD, the text mentions that the command may specify a drive (e.g., LOAD B:FOOBAR) on which to look for the .HEX file, but does not clearly spell out that LOAD will also place the resulting .COM file on the same drive.

Chapter 5, "Other Transient Programs and CP/M," is in three sections: utility programs, high level languages and application programs. The section on utilities deals with FORMAT, COPY, MOVCPM, and SYSGEN. Examples of formatting program names are listed, a typical run sequence is shown, and the process executed by the program is described in the most general terms. But, "You probably do not care what is done, or how, so long as the disk is properly prepared." The discussion of COPY is at a similar level.

The discussion of MOVCPM and SYSGEN is, to my mind, one of the least satisfactory parts of the book. The commands and the results and messages to be expected are well described, but it is all valueless because it is out of context. You are warned that "MOVCPM moves only the raw, bare bones of the CP/M operating system. Any additional drivers or other changes you have made to the BIOS section of CP/M (see Chapter 7) will not be moved." You are *not* warned that the Digital Research version of MOVCPM contains the BIOS for the Intel MDS (Microcomputer Development System) with its disk controller, and that few hardware vendors customize this program for their own controller. Nor is any mention made of the special area in the TPA where MOVCPM places the memory image of the relocated CP/M. Discussion of these programs belongs with discussion of the BIOS, or in a separate, very detailed and painstaking chapter on the relocation of CP/M. This is a tough subject, but any user may be faced with it, if only to accomodate application programs like BYE, which require protected space above CP/M. It needs far more extensive treatment than it gets in this book.

On the other hand, the section on high-level languages is very good indeed. It gives an overview of the characteristics of a number of languages for use on a CP/M system. There is useful information on what the code looks like, what the language can and cannot do for you, and some suggestions on how to pick a suitable language for your particular field of endeavor. The languages discussed include: interpreter and compiler Basic, with considerable detail on EBasic and CBasic; the C language; Cobol; Forth; Fortran; Pascal; and PL/I-80. "How to Run" paragraphs give the principal commands and options for specific implementations of these languages.

The section on application programs contains some general notes on the running of application programs, followed by brief descriptions of several word processors (Magic Wand, Scope, Electric Pencil, WordStar).

Chapter 6, "MP/M, CP/NET, and CP/M Derivatives," describes the purpose and functions of MP/M and CP/NET. The commands and functions peculiar to these operating systems are given in detail. In the discussion of CP/M Derivatives, the emphasis is on the differences between these and CP/M. As an ignoramus in regard to these other systems, I can say only that the introductory text and the details of the commands are clear and readable. Since I know Chapters 2 through 4 to be exceptionally accurate, I would trust the accuracy of this chapter also.

Chapter 7, "Technical Aspects of CP/M," is a very mixed bag. It begins by describing the general structure of CP/M, with both tabular and diagrammatic memory maps. These are followed by brief descriptions of the CCP and the BDOS, the cold start process, and a paragraph on how to use the BDOS functions in assembly language programs. The calling sequence is shown, and Table 7-1 lists all of the BDOS functions and their codes. So far, so good.

Now, however, we come to a section on modifying the BIOS. This seems to be predicated on two assumptions: 1) that the MOVCPM supplied by the vendor has been fully customized for your system; and 2) that you can successfully overlay your current BIOS with the modified version *in the operating area*. I am far from convinced that these assumptions are justified. There is no mention of the special area of the TPA where MOVCPM places the new memory image and where the SYSGEN expects to find it, nor of how to calculate the Read offset with which DDT can bring a BIOS to the right place in that area. There are plenty of warnings and suggestions which may work if you are modifying a system of the same size, but heaven help you if you try to overlay a modified BIOS on a relocated CP/M—this chapter certainly will not. And that, with the exception of a page on the IOBYTE and the codes for device names recognized by STAT, is it. There's nothing about directory entries, nothing about how CP/M allocates disk space, nothing technical at all.

Contrasting this skimpy and (I think) misleading material with the detailed content and high level of teaching competence found in Chapters 2 through 6, I found this chapter bitterly disappointing. But perhaps I am being unfair in asking my driving instructor to teach me maintenance as well.

Chapter 8, entitled "Putting It All Together: The Systems Approach": was missing from my proof copy, but will contain System Recommendations, Procedure Recommendations, and The Systems Approach.

Six appendices contain a CP/M Command Summary, ASCII character codes, Comparison of CP/M Versions 1.3, 1.4 and 2.0 in tabular form, CP/M Prompts, Diskette Selection (the size, type, and sectoring for various standard systems and controllers), and an extremely useful annotated bibliography. And, to restore one's faith in mankind, an excellent index. I chose about thirty items to spot-check, found index entries for them without difficulty, and the page references were accurate. There are positively *no* circular cross-references (Basilisk—see Weasel, Weasel—see Basilisk) which still occasionally creep into the works of publishers less careful than Osborne/McGraw-Hill. The book was written and edited with the aid of WordStar, and has been set in a very pleasing and legible typeface. ∎

A Disk Alignment Routine

by Lorin S. Mohler

This routine (written in the 'C' language,) is handy to use during disk drive alignment. It is for systems using the single-density Tarbell interface.

If you do your own disk drive alignment, you will find this program very handy.

Install the diskette containing ALIGN.COM in a working drive on the system. Execute ALIGN. Turn the disk drive power off and install the drive to be aligned. The drive must be configured as A, B, C or D. Turn the drive power on and insert an alignment diskette. Select the drive with the 'D' command. Select the track with the 'T' command. Load and unload the head with the 'L' and 'U' commands. Follow the recommended manufacturers alignment procedure for your drive. The head load or unload will be maintained at 'command:' only.

```
/*ALIGN. C

#define READ 0x8C
#define SEEK 0x12
#define BUSY 0x01
#define LOAD 0x0001
#define UNLOAD 0x0000
#define SECT 0xFA
#define DDATA 0xFB
#define DEXT 0xFC
#define WAIT 0xFC
#define DCOM 0xF8
#define DSTAT 0xF8

int head_load;                  /* global */

main()
{
char c;

head_load = UNLOAD;             /* initialize head condition */

  printf("rev. 9/23/80\n\n");
  while(1) {
      printf("command: ");      /* issue a prompt */

          while ( !kbhit())      /* maintain head */
              head();

          switch(c = toupper(getchar())) {    /* get command */

          case 'L':     printf("oad\n");
                        head_load = LOAD;
                        break;
          case 'U':     printf("nload\n");
                        head_load = UNLOAD;
                        break;
          case 'T':     printf("rack #: ");
                        seek();
                        break;
          case 'D':     printf("rive: ");
                        select();
                        break;
          case 'B':     printf("ye...\n");
                        exit();
          default:      printf(" ?\007");
                        pcmnds();


          }
      }
}
```

```
seek()                          /* move head to selected track */
{
  char c[80];
  int track;

  track = atoi(gets(c));
  if (track < 0 || track > 76) {     /* get track number */
                                     /* check legality */
      printf(" ?\007: range 0-76\n");
      return(-1);                    /* error flag */
  }
  outp(DDATA,track);                 /* set track */
  while (inp(DSTAT) & BUSY)
      ;
  outp(DCOM,SEEK);
  inp(WAIT);
  printf("\n");
  return(0);
}

select()                        /* select the drive */
{
  char drive;

  drive = toupper(getchar());        /* get drive */
  if (drive < 'A' || drive > 'D') {  /* check legality */
      printf(" ? \007: range A-D\n");
      return(-1);                    /* error flag */
  }
  outp(DEXT, ~(drive-'A') << 4 | 2); /* set drive */
  printf("\n");
}

head()                          /* maintain the head load or unload */
{
  if (head_load) {
          outp(SECT,1);         /* sector 1 */
          outp(DCOM,READ);      /* read loads the head */
  }
}

pcmnds()                        /* print the legal commands */
{
  printf (": (l)oad & (u)nload head, (t)rack, (d)rive & (b)ye\n");
}
```

This program was compiled using the BD Software C compiler written by Leor Zolman. The User's Guide that comes with the diskette is excellent. Also be sure to get a copy of *The C Programming Language* by Brian W. Kernighan and Dennis Ritchie (Prentice Hall 1978) for a proper description of the language. ■

Lorin S. Mohler, PO Box 8340, Anahiem, CA 92802

Bring the flavor of Unix To your Z80-based CP/M system with Unica

"Unicum: a thing unique in its kind, especially an example of writing. Unica: the plural of unicum."

The Unica: a unique collection of programs supporting many features of the Unix operating system never before available under CP/M. The Unica are more than software tools; they are finely crafted instruments of surgical quality. Some of the Unica are:

bc	-	binary file compare
cat	-	catenate files
cp	-	copy one or more files
dm	-	disk map and statistics
hc	-	horizontal file catenation
ln	-	create file links (aliases)
ls	-	directory lister
mv	-	move (rename) files, even across users
rm	-	remove files
sc	-	source file compare, with resynchronization
srt	-	in-memory file sorter
sr	-	search multiple files for a pattern
sp	-	spelling error detector, with 20,000 word dictionary

Each Unicum understands several flags ("options" or "switches") which control program alternatives. No special "shell" is needed; Unica commands are typed to the standard CP/M command interpreter. The Unica package supports several Unix-like facilities, like filename user numbers:

    sc data.bas;2 data.bas;3

(compares files belonging to user 2 and user 3); Wildcard patterns:

    rm *tmp* -v

(types each filename containing the letters TMP and asks whether to delete the file); I/O redirection:

    ls -a ►list

(writes a directory listing of all files to file "list"); Pipes:

    cat chap* ! sp ! srt ►lst:

(concatenates each file whose name starts with "chap", makes a list of mispelled words, sorts the list, and prints it on the listing device).

The Unica are written in XM-80, a low level language which combines rigorously checked procedure definition and invocation with the versatility of Z80 assembly language. XM-80 includes a language translator which turns XM-80 programs into source code for MACRO-80, the industry standard assembler from Microsoft. It also includes a MACRO-80 object library with over forty "software components", subroutine packages which are called to perform services such as piping, wildcard matching, output formatting, and device-independent I/O with buffers of any size from 1 to 64k bytes.

The source code for each Unicum main program (but not for the software component library) is provided. With the Unica and XM-80, you can customize each utility to your installation, and write your own applications quickly and efficiently. Programs which you write using XM-80 components are not subject to any licensing fee.

Extensive documentation includes tutorials, reference manuals, individual spec sheets for each component, and thorough descriptions of each Unicum.

Update policy: each Unica owner is informed when new Unica or components become available. At any time, and as often as you like, you can return the distribution disk with a $10 handling fee and get the current versions of the Unica and XM-80, with documentation for all new or changed software.

The Unica and XM-80 (which requires MACRO-80) are priced at $195, or $25 for the documentation. The Unica alone are supplied as *.COM executable files and are priced at $95 for the set, or $15 for the documentation. Software is distributed on 8" floppy disks for Z80 CP/M version 2 systems.

Knowlogy
"Shaping Knowledge for Evolving Worlds"
P.O. Box 283
Wilsonville, Oregon 97070

Visa/Mastercard customers call (503) 635-5701 after hours for next day shipment.

CP/M is a trademark of Digital Research; Unica is a trademark of Knowlogy; Unix is a trademark of Bell Telephone Labs; XM-80 is a trademark of Scientific Enterprises; Z80 is a trademark of Zilog Inc.

SOFTWARE DIRECTORY

Program Name: COMPRESS
Hardware System: CP/M
Minimum Memory Size: 16K
Language: 8080 Assembly
Description: COMPRESS is a group of four programs which perform data compression on ASCII data. There are two .COM files which will compress and decompress disk files, and two .REL files which are in subroutine form and act on memory use. Compression varies depending on the contents of the source file, but normally is in the range of 30% - 80%. Works on all 7 bit ASCII data.
Release: July 1981
Price: $50
Included with price: 8" CP/M disk with REL, COM, and DOC files.
Where to purchase it:
New Jersey Software Services
6 Village Circle
Westfield, NJ 07090

Program Name: SCREENMASTER
Hardware System: CP/M compatible. Any dumb terminal.
Minimum Memory Size: 48K recommended.
Language: CBasic-2.
Description: A complete screen handler facility, provided in source code for inclusion in programs requiring full/multi screen input. GOSUB SCREENMASTER. Upon return, multiple screens of validated input are available, in memory, for further use by the programmer. Programmer may insert CBasic-2 code at any of the many exits to

affect any degree of editing or control, overriding SCREENMASTER editing if desired. The programmer and, optionally, the terminal user have commands: GOTO (field) n, BACK/FORWARD n (fields), NEXT/PRIOR (screen), PRINT (screen), as well as SUBMIT and ABORT. Utilities are provided to create and test screens as well as to configure any dumb terminal.
Release: August 1981
Price: $195 Manual alone $25 Demo disk $10 additional.
Included with price: 90 page user manual, floppy disk with source code, demos and utilities.
Where to purchase it:
Marketing Essentials, Inc.
206 Mosher Ave.
Woodmere, NY 11598
(212) 580-3589

Program Name: FORTH by Timin Engineering, Release 2
Hardware System: 8080/8085/Z-80 CP/M or CDOS
Minimum Memory Size: 24K
Description: Enhanced version of FIG FORTH. A FORTH style editor with twenty commands is included, as well as virtual memory sub-system for disk I/O. The user may also make permanent additions to the resident FORTH vocabulary. A Z-80/8080 assembler is also included, allowing the user to create new FORTH definitions which compile directly into machine code. The IF...ELSE.., BEGIN...UNTIL, and BEGIN... WHILE... control structures may be included in assembler definitions. Other enhance-

ments include an interleaved disk for disk access. A 1024 byte disk block may be read or written in as little as 1/6 second. Eight of these blocks are maintained in RAM for immediate access and automatically swapped with others on the disk as they are needed.
Released: December 1980
Price: $95, IBM compatible 8" SD disk (other formats $110). Manual only: $20, applicable toward purchase on disk. California residents add 6% sales tax.
Included with price: CP/M compatible disk, user's manual and shipping by mail in U.S.
Authors: Dr. Mitchell E. Timin, and FIG
Where to pruchase:
Timin Engineering Co.
9575 Genesee Avenue, Suite E-2
San Diego, CA 92121
(714) 455-9008

Program Name: CONST
Hardware System: North Star & Apple
Language: Basic
Description: This program was written to do quantity and sizing take-offs for residential and small commercial structures. To operate the program, the user has only to answer questions concerning room sizes and type of construction.
Release: July 1981
Price: $75; listing only $60.
Included with price: Diskette, On-line Documentation, Support
Where to Purchase it:
Computing Interface
1918 Carnegie Lane #C
Redondo Beach, CA 90278

Software Directory, continued...

Program Name: Programmer's Apprentice
Hardware System: 8080/Z80, CP/M
Minimum Memory Size: 56K of RAM
Language: MBasic
Description: A program development that uses a macro-like language to define standard routines used by its code generator to create MBasic source code programs. Creates fully debugged programs in MBasic to provide screen prompted data input, data base management, file maintenance, and report generation. It is a visually oriented system designed to increase productivity and ease development of interactive application and report programs by performing most of the routine drudgery of programming.

To generate a program, first one creates the screen or report template on the CRT via the built-in screen editor. Then the definition modules use the screen/report template to define the various fields' attributes, eliminating most input errors, and selecting which fields will be the record control keys. Finally, the actual MBasic source code is generated.

Release: October 1981
Included with price: Manual, diskette, software including 80 column report generator.
Where to purchase it:
The Software Group
10471 S. Brookhurst
Anaheim, CA 92804
(714)535-5274

Program Name: WORD-C1, a text formatter
Hardware System: CP/M 2.2, an 8" dual diskette.
Minimum Memory Size: 60K
Language: Compiled
Description: A text formatter to prepare letters, memos, reports, manuals, documents and books. Commands set page length, line width, skip pages and text, indent, center text, etc. Line spacing, filing, adjusting, margin right justification and page numbering are automatically controlled. WORD has no limit to text size, no special hardware or modification. The Mail merge option lets you merge text with data files created by IDM-C1.

Released: September 1981
Price: $85
Included with price: 8" diskette, user's manual & postage.
Where to purchase it:
Micro Architect Inc.
96 Dothan Street
Arlington, MA 02174
(617)643-4713

Program Name: Disc-tionary
Hardware System: CP/M
Minimum Memory Size: 32K
Language: Z80 machine code
Description: High speed text proofreader. A thirty page document can be checked for errors in less than two minutes. Each unrecognized word can be added to the Disc-tionary, rejected, ignored etc. with a

DON'T MISS PHILADELPHIA AREA COMPUTER SHOW

Sponsored by Philadelphia Area Computer Society

PCAS-81

NOVEMBER 12-14
Philadelphia Sheraton Hotel
17th & J. F. Kennedy Boulevard

Featuring Applications of Computers in

- # EDUCATION
- # BUSINESS
- # LEISURE ACTIVITIES

PACS-81 covers a broad spectrum of interest among computer USERS. A special area of the Exhibit Hall will be reserved for

APPLICATIONS IN EDUCATION

and Computer Assisted Instruction. The most comprehensive educational exhibit to date.

BUSINESS AND HOBBY

application will also be exhibited. Commercial and Proprietary Software and Hardware will be on display including Special Exhibits on

COMPUTER MUSIC AND ROBOTICS.

CONFERENCE:
November 12-14 (Thurs., Fri., Sat.) Speakers, Workshops and Seminars.

EXHIBITS:
November 13 (Fri.) 11-9
November 14 (Sat.) 10-6

ADMISSION: $5

EXHIBITOR BOOTHS AVAILABLE
CONTACT:
HOWARD LUBERT
5127 Windward Lane
Bensalem, Pa. 19020
(215) 245-0554

Software Shops

—Arizona—

Creative Software Systems: Systems integration and custom software (BASIC, PASCAL, Z-80 assembler). Small business and word processing systems. 632 Camelot Dr., Sierra Vista, AZ 85635.
Phone (602)458-6063.

—California—

Clear Systems
309 Santa Monica Blvd., # 404
Santa Monica, CA 90401.
(213)394-7740.
(making complexity serve simplicity)

—Colorado—

Random Factors LTD.: Industrial test, control & data acquisition — Hi speed & accuracy for S100 & STD-BUS. From software to complete systems. W.K. Borsum, P.E., Random Factors LTD. Castle Rock, CO 80104. (303) 688-5338.
Nelson Engineering: We write applications software for all micro-based systems in Asembly language, Basic, and Pascal. (213) 390-2963; 13450 Maxella Ave. G185 Suite 142, Marina Del Rey, CA 90291.

—Massachusetts—

MICROFT INC.: Customization of CP/M-80, MP/M, CP/M-86 and other operating systems. Full range of consulting services in microsystems software (systems, utilities applications), product selection, hardware. Contact: Tom Campbell, Chief of Technical Staff, P.O. Box 128, E. Falmouth, MA 02536. Phone (617)563-3807.

—New Jersey—

New Jersey Software Services: Full range of CP/M, S-100 services.
— System design
    Business applications
    Real-time systems
    Mathematical analysis
— Software creation/customizing
    8080 assembly language
    Z-80 assembly language
    BASIC
    FORTRAN
— Product evaluation/selection
    Hardware Software
— In house training
— Telecommunication service
— Voice I/O applications
Contact C. A. Ryan, 6 Village Circle, Westfield, N.J. 07090 (201) 233-9297.

—Washington—

CHI ENERGY: Custom programs and package modification in Assembler, Basic & C languages; CP/M and real time systems. Contact: Mark A. Carlson, P.O. Box 55145, Seattle, WA 98155. (206)364-5463

single keystroke. Similar words may also be listed. After proofreading, the Disc-tionary leaves files containing the original unmarked text (BAK file), the marked text, and an alphabetized list of misspellings. As distributed, nearly 50,000 words can be recognized, and expansion allows up to four times that many. Many options including automatic suffix removal are available. All functions are performed by a single menu-driven program for ease of use. Two free "bug-fix" updates are included in the price.
Release: September 1981
Price: $79.00
Included with price: User manual, 8" diskette (manual available separately for $15.00).
Where to purchase it:
Stellarsoft Corporation
841 Blanchette Dr.
East Lansing, MI 48823
(517)332-2459

Program Name: Utilities Software Disk (DMM-1)
Hardware System: CP/M 2.x or MP/M system
Minimum Memory Size: 16K
Language: Object Code
Description: Disk contains the following programs:
    XDIR: Displays disk directory file names in alphabetic order and size of each file

name. Also shown are the number of bytes on disk, number of file names in use, space used, number of available file names and space. Works on single and double density floppy disks as well as with hard disks.
    EXTRACT: Will list a portion of a file between two label names. You do not have to list out whole file.
    STRIP: Removes hex code from a PRN file and turns it back into an ASM file.
    SORT: Creates symbol table from an assembly done with ASM that can be listed or used with Digital Research debugger SID.
    CONVERT: Changes all uncommented lower case characters to upper case. Handy for nice looking listings and for assemblers that will not accept lower case.
    STATUS: Provides system information such as memory available, TPA size, top of memory address, I/O assignments and more.
Release: September 1981
Price: $35 plus $1.50 shipping and handling
Included with price: Disk, 8" single density or 5-1/4" North Star
Where to purchase it:
Elliam Associates
2400 Bessemer St.
Woodland Hills, CA 91367
(213)348-4278 ∎

Get 12 issues of Creative Computing for the price of 8.

Some things are still cheaper by the dozen.

When you subscribe to *Creative Computing*, you get 12 issues for just $20. The same 12 issues would cost you $30 at the newsstand.

Why not enjoy *Creative Computing* all year long and save $10 at the same time.

To subscribe, call toll-free from 9 AM to 6 PM 800-631-8112. In New Jersey, call 201-540-0445. Or write to Creative Computing, Morris Plains, NJ 07950. We accept Visa, MasterCard and American Express.

Creative Computing is the leading magazine of small computer applications and software. It has in-depth reviews of new systems, peripherals and software. Also articles for both beginners and experts; columns about popular computers, programming techniques and new products; and complete program listings for your computer.

Alvin Toffler says, "I read *Creative Computing* not only for information about how to make the most of my own equipment but to keep an eye on how the whole field is emerging."

Why not join over 90,000 subscribers and save money at the same time? If you're clever enough to order a dozen.

**Multi-User Software Offers
Extra Performance And Features For
CP/M, MP/M, CP/Net-type Systems**

MuSYS Corporation, maker of multi-user microcomputer systems, is introducing MuDOS, a CP/M compatible operating system offering higher throughput, increased reliability and extra professional features for both single and multiple-user environments. The system may be customized to any Z80-based hardware configuration and used in place of CP/M, MP/M and CP/Net. Program loading under MuDOS is six times faster than CP/M. File processing functions average three to five times faster. MuDOS uses the extra registers and instructions available on the Z80 to speed processing of operating system calls.

Many standard features on MuDOS are either optional, extra cost or on available on CP/M. The most significant of these extra features are a buffer manager, a totally re-entrant file manager and an optional multiple print queuing capability. A sophisticated buffer manager reduces greatly the number of physical disk accesses required. The amount of buffering is a user parameter, and may be changed dynamically by the application program. A re-entrant file manager allows simultaneous accesses on different controllers, if the controller hardware supports DMA transfers. Also, higher priority tasks may interrupt long file operation. Optional automatic, concurrent print spooling, allows use of multiple print queues, forms types, fonts, and even hand-fed single sheet printing. This option satisfies a wide variety of data and word processing requirements.

MuDOS supports disk files up to 67 megabytes and drives up to 2000 megabytes and is usable with today's most advanced disk systems. Warm start and disk log-on displays are unnecessary. Allocation information is recorded on the disk. Disks may be changed any time updating is not in process. This eliminates disk system resets or warm starts which are particularly bothersome in other multi-user implementations.

MuDOS performs read-after-write verification of each disk update operation. Due to extensive buffering, this is done with little or no degradation in throughput. Any errors detected are reported with meaningful diagnostic messages, complete with alternatives for handling. This enhances reliability and permits graceful recovery from errors.

Modular architecture permits easy adaptation to different user environments. Module commonality ensures compatibility between different MuDOS versions. A relocating, linking, loader program determines system memory size, selects the configuration specified in a symbolic parameter file, and relocates, links, and loads the modules. Up to 30 modules are available in the resident portion of MuDOS, depending on system configurations elected. Each hardware dependent element is held within a separate relocatable module. Modules may be changed or replaced without massive reasasemblies or system generations.

Price: MuDOS (diskette and User's Guide) $300-$750 depending on configuration. For more information, contact Mr. Bill Schultz, MuSYS Corporation, 1451 E. Irvine Blvd., Suite 11, Tustin, CA 92680. (714)730-5692. TWX: 910-595-1967. Cable: MUSYSTSTN.

CompuPro Announces 64K Static RAM

CompuPro division of Godbout Electronics announces the RAM 17—a 64K fully static memory board for S-100/IEEE 696 computers. It is the highest density static RAM ever offered for the S-100/IEEE 696 bus on a standard 5 inch high board, and offers many advantages over its dynamic counterparts. Among them are faster access time guaranteed to run with 6 Mhz Z-80's and 10 Mhz 8086/88's, lower power dissipation (less than 2 watts), freedom from alpha particle soft-bit errors, and problem-free DMA operation.

It meets all the IEEE 696 specifications including 24-bit addressing, allowing up to 16 Mbytes of system memory, may be addressed on any 64K page boundary, and can be disabled in 16K blocks. The upper 8K block can also have 2K windows disabled to allow for memory-mapped peripherals. It uses the new 2K x 8-bit static RAM chips that are pin-compatible with 2716 type EPROM's, allowing up to 32K or EPROM to intermix with the RAM.

New Products, continued...

Available in Unkit, Assembled and Tested or Certified System Component high reliability program. Single quantity end-user prices are $1095, $1395 and $1595 respectively. Also available in a 48K configuration. CompuPro division of Godbout Electronics, P.O. Box 2355, Oakland Airport, CA 94614.

S-100 4-Port Serial I/O Interface

The CCS 2710 4-Port Serial I/O Interface permits independent programmable port control using four 8250 asynchronous communications elements with extensive control over I/O format and operation. Each port has full handshaking. The system interface

meets IEEE proposed S-100 standards. Word length, parity, modem, number of stop bits, and interrupt conditions are all software-determined. The baud rate may be programmed for any between 2K and 56K baud.

Driver routines and/or other programs may be stored on-board in a user-supplied 2K EPROM. The base address of the ROM is jumper-selectable at any 2K boundary. If the phantom output is jumper-enabled, the ROM routines memory overlay sharing the ROM's addresses.

For more information, contact: California Computer Systems, 250 Caribbean Drive, Sunnyvale, CA 94086, (408)734-5811.

Cromemco-Compatible 64K RAM Board

SUPERAM 4C is a 64K dynamic RAM board functionally equivalent to Cromemco's 64KZ RAM card. No user training or jumpers are required to achieve compatibility because the operation of all of the diagnostic LED's and user switches is identical. The board is organized in two blocks of 32K bytes. Each 32K block can be placed in any of eight different memory banks. Bank selection is available on 32K byte boundaries. Address and bank assignment of each 32K block is selectable via simple switch settings. Up to eight boards can be used with bank select to expand memory from 64K to 512K. Memory refresh is totally transparent to the processor. Memory access time is 250 nsec. Operating speed is 4MHz with no wait states. The price of the board is $995. Piiceon Inc., 2350 Bering Drive, San Jose, CA 95131; (408)946-8030.